

# Minimum Spectral Connectivity Projection Pursuit For Unsupervised Classification

David P. Hofmeyr<sup>1</sup>

Nicos G. Pavlidis<sup>2</sup>

Idris A. Eckley<sup>3</sup>

<sup>1</sup> STOR-i center for doctoral training, Lancaster University

<sup>2</sup> Department of management science, Lancaster University

<sup>3</sup> Department of mathematics and statistics, Lancaster University

## Abstract

We study the problem of determining the optimal low dimensional projection for maximising the separability of a binary partition of an unlabelled dataset, as measured by spectral graph theory. This is achieved by finding projections which minimise the second eigenvalue of the Laplacian matrices of the projected data, which corresponds to a non-convex, non-smooth optimisation problem. We show that the optimal univariate projection based on spectral connectivity converges to the vector normal to the maximum margin hyperplane through the data, as the scaling parameter is reduced to zero. This establishes a connection between connectivity as measured by spectral graph theory and maximal Euclidean separation. It also allows us to apply our methodology to the problem of finding large margin linear separators. The computational cost associated with each eigen-problem is quadratic in the number of data. To mitigate this problem, we propose an approximation method using microclusters with provable approximation error bounds. We evaluate the performance of the proposed method on a large collection of benchmark datasets and find that it compares favourably with existing methods for projection pursuit and dimension reduction for unsupervised data partitioning.

## Keywords

Spectral clustering, dimension reduction, projection pursuit, eigenvalue optimisation, maximum margin clustering

## 1 Introduction

The classification of unlabelled data is fundamental to many statistical and machine learning applications. Such applications arise in the context of clustering and semi-supervised classification. Underpinning these tasks is the assumption of a clusterable structure within the data, and importantly that this structure is relevant to the classification task. The assumption of a clusterable structure, however, begs the question of how a cluster should be defined. Centroid based methods, such as the ubiquitous  $k$ -means algorithm, define clusters in reference to single points, or centers (Leisch, 2006). In the non-parametric statistical approach to clustering, clusters are associated with the modes of a probability density function from which the data are assumed to arise (Hartigan, 1975, Chapter 11). We consider the definition as given in the context of graph partitioning, and the relaxation given by spectral clustering. Spectral clustering has gained considerable interest in recent years due to its strong performance in diverse application areas. In this context clusters are defined as strongly connected components of a graph defined over the data, wherein vertices correspond to data points and edge weights represent pairwise similarities (von Luxburg, 2007).

The minimum cut graph problem seeks to partition a graph such that the sum of the edges connecting different components of the partition is minimised. To avoid partitions containing small sets of vertices, a normalisation is introduced which helps to emphasise more balanced partitions. The normalisation, however, makes the problem NP-hard (Wagner and Wagner, 1993), and so a continuous relaxation is solved instead. The relaxed problem, known as spectral clustering, is solved by the eigenvectors of the *graph Laplacian* matrices. We give a brief introduction to spectral clustering in Section 3.

Crucial to all cluster definitions is the relevance of spatial similarity of points. In multivariate data analysis, however, the presence of irrelevant or noisy features can significantly obscure the spatial structure in a data set. Moreover, in very high dimensional applications the curse of dimensionality can make spatial similarities unreliable for distinguishing clusters (Steinbach et al., 2004; Beyer et al., 1999). Dimension reduction techniques seek to mitigate the effect of irrelevant features and of the curse of dimensionality by finding low dimensional representations of a set of data which retain as much information as possible. Most commonly these low dimensional representations are defined by the projection of the data into a linear subspace. Information retention is crucial for the success of any subsequent tasks. For unsupervised classification this information must, therefore, be relevant in the context of cluster structure. Classical dimension reduction techniques such as principal component analysis (PCA) cannot guarantee the structural relevance of the low dimensional subspace. Moreover a single subspace may not suffice to distinguish all clusters, which may have their structures defined within differing subspaces. Recently a number of dimension reduction methods with an explicit objective which is relevant to cluster structure have been proposed (Krause and Liebscher, 2005; Niu et al., 2011; Pavlidis et al., 2016). We discuss these briefly in Section 2.

We consider the problem of learning the optimal subspace for the purpose of data bi-partitioning, where optimality is measured by the connectivity of the projected data, as defined in spectral graph theory. We formulate the problem in the context of *projection pursuit*; a class of optimisation problems which aim to find *interesting* subspaces within potentially high dimensional data sets, where interestingness is captured by a predefined objective, called the *projection index*. With very few exceptions, the optimisation of the projection index does not admit a closed form solution, and is instead numerically optimised. The projection indices considered in the proposed method are the second smallest eigenvalues of the graph Laplacians, which measure the quality of a binary partition arising from the normalised minimum cut graph problem. These eigenvalues are non-smooth and non-convex, and so specialised techniques are required to optimise them. We establish conditions under which they are Lipschitz and almost everywhere continuously differentiable, and discuss how to find local optima with guaranteed convergence properties.

In this paper we establish an asymptotic connection between optimal univariate subspaces for bi-partitioning based on spectral graph theory, and maximum margin hyperplanes. Formally, we show that as the scaling parameter defining pairwise similarities is reduced to zero, the optimal univariate subspace for bi-partitioning converges to the subspace normal to the largest margin hyperplane through the data. This establishes a theoretical connection between connectivity as measured by spectral graph theory and maximal Euclidean separation. It also provides an alternative methodology for learning maximum margin clustering models, which have attracted considerable interest in recent years (Xu et al., 2004; Zhang et al., 2009). We introduce a way of modifying the similarity function which avoids focusing on outliers, and allows us to further control the balance of the induced partition. The importance of controlling this balance has been observed in the context of large margin clustering (Xu et al., 2004; Zhang et al., 2009) and low density separators (Pavlidis et al., 2016).

The computation cost associated with the eigen-problem underlying our projection index is quadratic in the number of data. To mitigate this computational burden we propose a data preprocessing step using micro-clusters which significantly speeds up the optimisation. We establish theoretical error bounds for this approximation method, and provide a sensitivity study which shows no degradation in clustering performance, even for a coarse approximation.

The remainder of the paper is organised as follows. In Section 2 we briefly discuss related work on dimension reduction for unsupervised data partitioning. A brief outline of spectral clustering is provided in Section 3. Section 4 presents the methodology for finding optimal projections to perform binary partitions. Section 5 describes the theoretical connection between optimal subspaces for spectral bi-partitioning and maximum margin hyperplanes. In Section 6 we discuss an approximation method in which the computational speed associated with finding the optimal subspace can be significantly improved, with provable approximation error bounds. Experimental results and sensitivity analyses are presented in Section 7, while Section 8 is devoted to concluding remarks.

## 2 Related Work

The literature on clustering high dimensional data is vast, and we will focus only on methods with an explicit dimension reduction formulation, as in projection pursuit. Implicit dimension reduction methods based on learning sparse covariance matrices (which impose an implicit low dimensional projection of the data/clusters), such as quadratic discriminant analysis, can be limited by the assumption that clusters are determined by their covariance matrices. Projection pursuit approaches can be made more versatile by defining objectives which admit more general cluster definitions.

Principal component analysis and independent component analysis have been used in the context of clustering, however their objectives do not correspond exactly with those of the clustering task and the justification of their use is based more on common-sense reasoning. Nonetheless, these methods have shown good empirical performance on a number of applications (Boley, 1998; Tasoulis et al., 2010; Kriegel et al., 2009). Some recent approaches to projection pursuit for clustering rely on the non-parametric statistical notion clusters, i.e., that clusters are regions of high density in a probability distribution from which the data are assumed to have arisen. Krause and Liebscher (2005) proposed using as projection index the *dip statistic* (Hartigan and Hartigan, 1985) of the projected data. The dip is a measure of departure from unimodality, and so maximising the dip tends to projections which have multimodal marginal density, and therefore separate high density clusters. The authors establish that the dip is differentiable for any projection vector onto which the projected data are unique, and use a simple gradient ascent method to find local optima.

The minimum density hyperplane approach (Pavlidis et al., 2016) is posed as a projection pursuit for the univariate subspace normal to the hyperplane with minimal integrated density along it, thereby establishing regions of low density which separate the modes of the underlying probability density. The projection index in this case is the minimum of the kernel density estimate of the projected data, penalised to avoid hyperplanes which do not usefully split the data. The authors show an asymptotic connection between the hyperplane with minimal integrated density and the maximum margin hyperplane. The result we show in Section 5 therefore establishes that the optimal subspace for bi-partitioning based on spectral connectivity is asymptotically connected with the minimum integrated density hyperplane.

A number of direct approaches to maximum margin clustering have also been proposed (Xu et al., 2004; Zhang et al., 2009). These can be viewed as a projection pursuit for the subspace normal to the maximum margin hyperplane intersecting the data. The iterative support vector regression approach (Zhang et al., 2009) uses support vector methods and so for the linear kernel explicitly learns the corresponding projection vector,  $v$ .

Most similar to our work is that of Niu et al. (2011), who also proposed a method for dimension reduction based on spectral clustering. The authors show an interesting connection between optimal subspaces for spectral clustering and *sufficient dimension reduction*. For the case of a binary partition, their objective is equivalent to one of the objectives we consider, i.e., that of minimising the second smallest eigenvalue of the normalised Laplacian (cf. Sections 3 and 4). However, our methodology differs substantially from theirs. Niu et al. (2011) define their objective by

$$\max_{U, W} \quad \text{trace}(U^\top D^{-1/2} A D^{-1/2} U) \quad (1a)$$

$$s.t. \quad U^\top U = I \quad (1b)$$

$$A_{ij} = s(\|W^\top x_i - W^\top x_j\|) \quad (1c)$$

$$W^\top W = I. \quad (1d)$$

The matrix  $A$  is the affinity matrix containing pairwise similarities of points projected into the subspace defined by  $W$ , and  $D$  is the diagonal degree matrix of  $A$ , with  $i$ -th diagonal element equal to the  $i$ -th row sum of  $A$ . Further details of these objects can be found in Section 3. The approach used by the authors to maximise this objective alternates between using spectral clustering to determine the columns of  $U$ , and then using a gradient ascent method to maximise  $\text{trace}(U^\top D^{-1/2} A D^{-1/2} U)$

over  $W$ , where the dependence of this objective on the projection matrix  $W$  is through Eq. (1c). Within this gradient ascent step the matrices  $U$  and  $D$  are kept fixed. This process is iterated until convergence. However, the authors do not address the fact that the matrix  $D$  is determined by  $A$ , and therefore depends on the projection matrix  $W$ . An ascent direction for the objective assuming a fixed  $D$  is therefore not necessarily an ascent direction for the overall objective. Despite this fact the method has shown good empirical performance on a number of problems (Niu et al., 2011). In Section 4 we derive expressions for the gradient of the overall objective, which allows us to optimise it directly.

### 3 Background on Spectral Clustering

In this section we provide a brief introduction to spectral clustering, with particular attention to bi-partitioning, which underlies the focus of this work. Bi-partitioning using spectral clustering has been considered previously by Shi and Malik (2000), where a full clustering can be obtained by recursively inducing bi-partitions of (subsets of) the data. With a data sample,  $X = \{x_1, \dots, x_N\}$ , spectral clustering associates a graph  $G = (V, E)$ , in which vertices correspond to observations, and the *undirected* edges assume weights equal to the pairwise *similarity* between observations. Pairwise similarities can be determined in a number of ways, including nearest neighbours and similarity metrics. In general, similarities are determined by the spatial relationships between points, and pairs which are closer are assigned higher similarity than those which are more distant.

The information in  $G$  can be represented by the *adjacency*, or affinity matrix,  $A \in \mathbb{R}^{N \times N}$ , with  $A_{ij} = E_{ij}$ . The *degree* of each vertex  $v_i$  is defined as,  $d_i = \sum_{j=1}^N A_{ij}$ . The *degree matrix*,  $D$ , is then defined as the diagonal matrix with  $i$ -th diagonal element equal to  $d_i$ . For a subset  $C \subset X$ , the size of  $C$  can be defined either by the cardinality of  $C$ ,  $|C|$ , or by the *volume* of  $C$ ,  $\text{vol}(C) = \sum_{i: x_i \in C} d_i$ .

**Definition** The *normalised min-cut graph problem* for a binary partition is defined as the optimisation problem

$$\min_{C \subset X} \sum_{i,j: x_i \in C, x_j \in X \setminus C} A_{ij} \left( \frac{1}{\text{size}(C)} + \frac{1}{\text{size}(X \setminus C)} \right). \quad (2)$$

It has been shown (Hagen and Kahng, 1992; Shi and Malik, 2000) that the two normalised min-cut graph problems (corresponding to the two definitions of size) can be formulated in terms of the *graph Laplacian* matrices,

$$(\text{standard}) \ L = D - A, \quad (3)$$

$$(\text{normalised}) \ L_{\text{norm}} = D^{-1/2} L D^{-1/2}, \quad (4)$$

as follows. For  $C \subset X$  define  $u^C \in \mathbb{R}^N$  to be the vector with  $i$ -th entry,

$$u_i^C = \begin{cases} \sqrt{\text{size}(X \setminus C) / \text{size}(C)}, & \text{if } x_i \in C \\ -\sqrt{\text{size}(C) / \text{size}(X \setminus C)}, & \text{if } x_i \in X \setminus C. \end{cases} \quad (5)$$

For  $\text{size}(C) = |C|$ , the optimisation problem in (2) can be written as,

$$\min_{C \subset X} (u^C)^\top L u^C \quad \text{s.t.} \quad u^C \perp \mathbf{1}, \ \|u^C\| = \sqrt{N}. \quad (6)$$

Similarly, if  $\text{size}(C) = \text{vol}(C)$  the problem in (2) is equivalent to,

$$\min_{C \subset X} (u^C)^\top L u^C \quad \text{s.t.} \quad D u^C \perp \mathbf{1}, \ (u^C)^\top D u^C = \text{vol}(V). \quad (7)$$

Both problems in (6) and (7) are NP-hard (Wagner and Wagner, 1993), and so continuous relaxations of these, in which the discreteness condition on  $u^C$  given in Eq. (5) is removed, are solved instead (Hagen and Kahng, 1992; Shi and Malik, 2000). The solutions to the relaxed problems are given by the

second eigenvector of  $L$  and the second eigenvector of the generalised eigen equation  $Lu = \lambda Du$  respectively, the latter thus equivalently solved by  $D^{-1/2}u$ , where  $u$  is the second eigenvector of  $L_{\text{norm}}$ . In particular, we have

$$\lambda_2(L) \leq \frac{1}{N}(u^S)^\top Lu^S \quad (8)$$

$$\lambda_2(L_{\text{norm}}) \leq \frac{1}{\text{vol}(V)}(u^N)^\top Lu^N, \quad (9)$$

where  $\lambda_2(L)$  and  $\lambda_2(L_{\text{norm}})$  are the second eigenvalues of  $L$  and  $L_{\text{norm}}$  and  $u^S$  and  $u^N$  are the solutions to (6) and (7) respectively.

The following properties of the matrices  $L$  and  $L_{\text{norm}}$  will be useful in establishing our proposed methodology and the associated theoretical results. These properties can be found in (von Luxburg, 2007, Propositions 2 and 3).

1. For any  $v \in \mathbb{R}^N$  we have

$$v^\top Lv = \frac{1}{2} \sum_{i,j} A_{ij}(v_i - v_j)^2 \quad (10)$$

$$v^\top L_{\text{norm}}v = \frac{1}{2} \sum_{i,j} A_{ij} \left( \frac{v_i}{\sqrt{d_i}} - \frac{v_j}{\sqrt{d_j}} \right)^2. \quad (11)$$

2.  $L$  and  $L_{\text{norm}}$  are symmetric and positive semi-definite.
3. The smallest eigenvalue of  $L$  is 0 with corresponding eigenvector  $\mathbf{1}$ , the constant 1 vector
4. The smallest eigenvalue of  $L_{\text{norm}}$  is 0 with corresponding eigenvector  $D^{1/2}\mathbf{1}$ .

The extension of clustering via the normalised min-cut to a  $K$ -partition of the data is similar, and can be solved approximately by the first  $K$  eigenvectors of either  $L$  or  $L_{\text{norm}}$  (von Luxburg, 2007).

## 4 Projection Pursuit for Spectral Connectivity

In this section we study the problem of minimising the second eigenvalue of the graph Laplacian matrices of the projected data. If the projected data are split in two through spectral clustering, then the projection that minimises the second eigenvalue of the corresponding graph Laplacian minimises the connectivity of the two components, as measured by spectral graph theory. Note that while we discuss explicitly the minimisation of the second eigenvalue, the methodology we present in fact applies to an arbitrary eigenvalue of the graph Laplacians. As a result, the method discussed herein trivially extends to the problem of determining a  $K$ -partition by minimising the sum of the  $K$  smallest eigenvalues of the Laplacians.

To begin with, let  $X = \{x_1, \dots, x_N\}$  be a  $d$ -dimensional data set and let  $V \in \mathbb{R}^{d \times l}$  be a *projection matrix*, where  $l \in \mathbb{N}$  is the dimension of the projection, and the columns of  $V$ ,  $\{V_1, \dots, V_l\}$ , have unit norm. With this formulation it is convenient to consider a parameterisation of  $V$  through polar coordinates as follows. Let  $\Theta = [0, \pi)^{(d-1) \times l}$  and for  $\theta \in \Theta$ , the projection matrix  $V(\theta) \in \mathbb{R}^{d \times l}$  is given by,

$$V(\theta)_{ij} = \begin{cases} \cos(\theta_{ij}) \prod_{k=1}^{i-1} \sin(\theta_{kj}), & i = 1, \dots, d-1 \\ \prod_{k=1}^{d-1} \sin(\theta_{kj}), & i = d. \end{cases} \quad (12)$$

From this we define the  $l$  dimensional *projected data set* by  $P(\theta) = \{p(\theta)_1, \dots, p(\theta)_N\} = \{V(\theta)^\top x_1, \dots, V(\theta)^\top x_N\}$ , and we let  $L(\theta)$  (resp.  $L_{\text{norm}}(\theta)$ ) be the Laplacian (resp. normalised Laplacian) of the graph of  $P(\theta)$ . Edge weights are determined by a positive function  $s : (\mathbb{R}^l)^N \times \{1 \dots N\}^2 \rightarrow \mathbb{R}^+$ , in

that the affinity matrix is given by  $A(\boldsymbol{\theta})_{ij} := s(P(\boldsymbol{\theta}), i, j)$ . In the simplest case we may imagine  $s$  being fully determined by the Euclidean distance between two elements of the projected data, i.e.,  $s(P(\boldsymbol{\theta}), i, j) = k(\|p(\boldsymbol{\theta})_i - p(\boldsymbol{\theta})_j\|)$ , for some function  $k : \mathbb{R} \rightarrow \mathbb{R}^+$ . However we prefer to allow for a more general definition. We discuss this further in Section 4.3.

Henceforth we will use  $\lambda_i(\cdot)$  to be the  $i$ -th (smallest) eigenvalue of its (in all cases herein) real symmetric matrix argument. The objectives  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  are, in general, non-convex and non-smooth in  $\boldsymbol{\theta}$ , and so specialised techniques are required to optimise them. In the following subsections we investigate their differentiability properties, and discuss how alternating between a naive gradient descent method and a descent step based on a directional derivative can be used to find locally optimal solutions.

## 4.1 Continuity and Differentiability

In this subsection we explore the continuity and differentiability properties of the second eigenvalue of the graph Laplacians, viewed as a function of the *projection angle*,  $\boldsymbol{\theta}$ . We will view the data set  $X$  as a  $d \times N$  matrix with  $i$ -th column equal to  $x_i$ , and similarly the projected data set as an  $l \times N$  matrix,  $P(\boldsymbol{\theta}) = V(\boldsymbol{\theta})^\top X$ , with  $i$ -th column  $p(\boldsymbol{\theta})_i$ .

**Lemma 1** *Let  $X = \{x_1, \dots, x_N\} \subset \mathbb{R}^d$  and let  $s(P, i, j)$  be Lipschitz continuous in  $P \in \mathbb{R}^{l \times N}$  for fixed  $i, j \in \{1 \dots N\}$ . Then  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  are Lipschitz continuous in  $\boldsymbol{\theta}$ .*

**Proof** We show the case of  $L(\boldsymbol{\theta})$ , where that of  $L_{\text{norm}}(\boldsymbol{\theta})$  is similar. The result follows from the fact that  $L(\boldsymbol{\theta})$  is element-wise Lipschitz as a composition of Lipschitz functions ( $V(\boldsymbol{\theta})$  is Lipschitz in  $\boldsymbol{\theta}$  as a collection of products of Lipschitz functions) and the fact that

$$|\lambda_i(L(\boldsymbol{\theta})) - \lambda_i(L(\boldsymbol{\theta}'))| \leq \|L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}')\| \leq N \sqrt{\max_{ij} |L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}')|_{ij}},$$

where the first inequality is due to Weyl (1912), and the second comes from Schur's inequality (Schur, 1911).

Rademacher's theorem therefore establishes that both objectives are almost everywhere differentiable (Polak, 1987). This almost everywhere differentiability can also be seen by considering that simple eigenvalues of real symmetric matrices are differentiable, e.g. Magnus (1985), and establishing that under certain conditions on the function  $s$  the eigenvalues of  $L(\boldsymbol{\theta})$  and  $L_{\text{norm}}(\boldsymbol{\theta})$  are simple for almost all  $\boldsymbol{\theta}$ .

Tao and Vu (2014) have shown that the real symmetric matrices with non-simple spectrum lie in a subspace of co-dimension 2. If we denote the space of real valued  $N \times N$  symmetric matrices by  $\mathcal{S}_N$ , and denote this subspace by  $S$ , then  $\mathcal{S}_N \setminus S$  is open and dense in  $\mathcal{S}_N$ . Sufficient conditions on the function  $s$  for the almost everywhere simplicity of  $\lambda_2(L(\boldsymbol{\theta}))$  (resp.  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$ ) are therefore that it is continuous in  $P$  for each  $i, j$  and for all  $B \in \mathcal{S}_N$  and  $U$  open in  $\Theta$ ,  $\exists \boldsymbol{\theta} \in U$  s.t.  $\text{trace}(L(\boldsymbol{\theta})B) \neq 0$  (resp.  $\text{trace}(L_{\text{norm}}(\boldsymbol{\theta})B) \neq 0$ ). Continuity of  $s$  ensures continuity of the functions  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$ , and therefore the openness of the preimage of  $\mathcal{S}_N \setminus S$ . The latter condition ensures that for each open  $U \subset \Theta$ , the span of the image of  $U$  under  $\lambda_2(\cdot)$  is  $\mathcal{S}_N$ . Therefore, in every open  $U \subset \Theta$ ,  $\exists \boldsymbol{\theta} \in U$  s.t.  $L(\boldsymbol{\theta}) \notin S$ . Therefore the pre-image of  $\mathcal{S}_N \setminus S$  is dense in  $\Theta$ .

Generalised gradient based optimisation methods are the natural framework for finding the optimal subspace for spectral bi-partitioning. Eigenvalue optimisation is, in general, a challenging problem due to the fact that eigenvalues are not differentiable where they coincide. The majority of approaches in the literature focus on the problems of minimising the largest eigenvalue or the sum of a predetermined number of largest eigenvalues (Overton and Womersley, 1993). Both of these problems tend to lead to a coalescence of eigenvalues, making the issue of non-differentiability especially problematic. Conversely

the minimisation of the smallest eigenvalue tends to lead to a separation of eigenvalues, and so non-differentiability is less of a concern (Lewis and Overton, 1996).

If the similarity function  $s$  is strictly positive, then  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  are bounded away from zero, and hence minimising these has the same benefits as does minimising the smallest eigenvalue in general, in that the corresponding optimisation tends to separate them from other eigenvalues. Despite this practical advantage, the simplicity of  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  is not guaranteed over the entire optimisation. We discuss a way of handling points of non-differentiability in Section 4.2. This approach uses the directional derivative formulation given by Overton and Womersley (1993), and allows us to find descent directions which also tend to lead to a decoupling of eigenvalues.

Global convergence of gradient based optimisation algorithms relies on the continuity of the derivatives (where they exist). To establish this continuity, we first derive expressions for the derivatives of  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  as a function of  $\boldsymbol{\theta}$ . Theorem 1 of Magnus (1985) provides a useful formulation of eigenvalue derivatives. If  $\lambda$  is a simple eigenvalue of a real symmetric matrix  $M$ , then  $\lambda$  is infinitely differentiable on a neighbourhood of  $M$ , and the differential at  $M$  is given by

$$d\lambda = u^\top d(M)u, \quad (13)$$

where  $u$  is the corresponding eigenvector. Let us assume that  $s(P, i, j)$  is differentiable in  $P \in \mathbb{R}^{l \times N}$  for fixed  $i, j \in \{1 \dots N\}$ . For brevity we temporarily drop the notational dependence on  $\boldsymbol{\theta}$  and denote the second eigenvalue of the Laplacian by  $\lambda$ , and the corresponding eigenvector by  $u$ . The derivative  $D_{\boldsymbol{\theta}}\lambda$  is given by the  $(d-1) \times l$  matrix with  $i$ -th column  $D_{\boldsymbol{\theta}_i}\lambda$ , where we consider the chain rule decomposition  $D_{\boldsymbol{\theta}_i}\lambda = D_P \lambda D_V P D_{\boldsymbol{\theta}_i} V$ . Here  $D \cdot$  is the differential operator. Since only the  $i$ -th column of  $V$  depends on  $\boldsymbol{\theta}_i$ , and only the  $i$ -th row of  $P$  depends on  $V_i$ , this product can be simplified as  $D_{\boldsymbol{\theta}_i}\lambda = D_{P_i} \lambda D_{V_i} P_i D_{\boldsymbol{\theta}_i} V_i$ , where  $P_i$  is used to denote the  $i$ -th row of  $P$ , while  $V_i$  and  $\boldsymbol{\theta}_i$  are, as usual, the  $i$ -th columns of  $V$  and  $\boldsymbol{\theta}$  respectively. We provide expressions for each of these terms below.

We first consider the standard Laplacian  $L$ . By Eq. (13) we have  $d\lambda = u^\top d(L)u = u^\top d(D)u - u^\top d(A)u$ . Now,

$$\frac{\partial D_{ii}}{\partial P_{mn}} = \sum_{j=1}^N \frac{\partial A_{ij}}{\partial P_{mn}} = \sum_{j=1}^N \frac{\partial s(P, i, j)}{\partial P_{mn}}, \text{ and } \frac{\partial A_{ij}}{\partial P_{mn}} = \frac{\partial s(P, i, j)}{\partial P_{mn}}, \quad (14)$$

and so,

$$\frac{\partial \lambda}{\partial P_{mn}} = u^\top \frac{\partial L}{\partial P_{mn}} u = \frac{1}{2} \sum_{i,j} (u_i - u_j)^2 \frac{\partial s(P, i, j)}{\partial P_{mn}}. \quad (15)$$

For the normalised Laplacian,  $L_{\text{norm}}$ , consider first

$$\begin{aligned} d(L_{\text{norm}}) &= d(D^{-1/2} L D^{-1/2}) \\ &= d(D^{-1/2}) L D^{-1/2} + D^{-1/2} d(D) D^{-1/2} - D^{-1/2} d(A) D^{-1/2} + D^{-1/2} L d(D^{-1/2}). \end{aligned}$$

We will again use  $\lambda$  and  $u$  to denote the second eigenvalue and corresponding eigenvector. Using the fact that  $L D^{-1/2} u = \lambda D^{1/2} u$ ,

$$\begin{aligned} d\lambda &= u^\top d(D^{-1/2}) L D^{-1/2} u + u^\top D^{-1/2} d(D) D^{-1/2} u - u^\top D^{-1/2} d(A) D^{-1/2} u \\ &\quad + u^\top D^{-1/2} L d(D^{-1/2}) u \\ &= \lambda u^\top d(D^{-1/2}) D^{1/2} u + u^\top D^{-1/2} d(D) D^{-1/2} u - u^\top D^{-1/2} d(A) D^{-1/2} u \\ &\quad + \lambda u^\top D^{1/2} d(D^{-1/2}) u \\ &= \lambda u^\top d(I) u + (1 - \lambda) u^\top D^{-1/2} d(D) D^{-1/2} u - u^\top D^{-1/2} d(A) D^{-1/2} u, \end{aligned}$$

since  $d(D^{-1/2}) D D^{-1/2} + D^{-1/2} d(D) D^{-1/2} + D^{-1/2} D d(D^{-1/2}) = d(D^{-1/2} D D^{-1/2}) = d(I) = \mathbf{0}$ , we have,

$$\begin{aligned} d\lambda &= (1 - \lambda) u^\top D^{-1/2} d(D) D^{-1/2} u - u^\top D^{-1/2} d(A) D^{-1/2} u \\ &= u^\top D^{-1/2} d(L) D^{-1/2} u - \lambda u^\top D^{-1/2} d(D) D^{-1/2} u. \end{aligned}$$

Therefore,

$$\frac{\partial \lambda}{\partial P_{mn}} = \frac{1}{2} \sum_{i,j} \left( \frac{u_i}{\sqrt{d_i}} - \frac{u_j}{\sqrt{d_j}} \right)^2 \frac{\partial s(P, i, j)}{\partial P_{mn}} - \lambda \sum_{i,j} \frac{u_i^2}{d_i} \frac{\partial s(P, i, j)}{\partial P_{mn}}. \quad (16)$$

The component  $D_{V_i} P_i$  is simply the  $N \times d$  transposed data matrix, and the  $d \times (d-1)$  matrix,  $D_{\theta_i} V_i$ , is given by

$$D_{\theta_i} V_i = \begin{bmatrix} -\sin(\theta_{1i}) & 0 & \dots & 0 \\ \cos(\theta_{1i}) \cos(\theta_{2i}) & -\sin(\theta_{1i}) \sin(\theta_{2i}) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \cos(\theta_{1i}) \prod_{k=2}^{d-1} \sin(\theta_{ki}) & \cos(\theta_{2i}) \prod_{k \neq 2} \sin(\theta_{ki}) & \dots & \cos(\theta_{d-1,i}) \prod_{k=1}^{d-2} \sin(\theta_{ki}) \end{bmatrix}. \quad (17)$$

Having derived expressions for the derivatives of  $\lambda_2(L(\theta))$  and  $\lambda_2(L_{\text{norm}}(\theta))$ , we can address their continuity properties. The components  $D_V P D_{\theta} V$  clearly form a continuous product in  $\theta$ . The continuity of the elements  $\partial \lambda / \partial P_{mn}$  can be reduced to addressing the continuity of the eigenvalue itself, of its associated eigenvector and a continuity assumption on the derivative of the function  $s$ . It is well known that the eigenvalues of a matrix are continuous, while the continuity of the elements of the eigenvector come from the fact that we have assumed  $\lambda$  to be simple (Magnus, 1985). We provide full expressions for the derivatives of  $\lambda_2(L(\theta))$  and  $\lambda_2(L_{\text{norm}}(\theta))$ , for the similarity function used in our experiments, in Appendix A.

The eigenvalues of a real symmetric matrix can be expressed as the difference between two convex matrix functions (Fan, 1949). If the similarity function,  $s$ , is Lipschitz continuous and differentiable we therefore have that  $\lambda_2(L(\theta))$  and  $\lambda_2(L_{\text{norm}}(\theta))$  are directionally differentiable everywhere. Overton and Womersley (1993) describe a way of expressing the directional derivative of the sum of the  $k$  largest eigenvalues of a matrix whose elements are continuous functions of a parameter, at a point of non-simplicity of the  $k$ -th largest eigenvalue. We will discuss the case of  $\lambda_2(L(\theta))$ , where  $\lambda_2(L_{\text{norm}}(\theta))$  is analogous. If we denote the sum of the  $k$  largest eigenvalues of  $L(\theta)$  by  $F^k(\theta)$  then,

$$\lambda_2(L(\theta)) = F^{N-1}(\theta) - F^{N-2}(\theta). \quad (18)$$

Now suppose that  $\theta$  is such that

$$\begin{aligned} \lambda_N(L(\theta)) &\geq \dots \geq \lambda_{N-r+1}(L(\theta)) > \\ \lambda_{N-r}(L(\theta)) &= \dots = \lambda_{N-k+1}(L(\theta)) = \dots = \lambda_{N-r-t+1}(L(\theta)) > \\ \lambda_{N-r-t}(L(\theta)) &\geq \dots \geq \lambda_1(L(\theta)). \end{aligned}$$

That is, the  $k$ -th largest eigenvalue has multiplicity  $t$  and  $k-r$  are included in the sum defining  $F^k(\theta)$ . Then the directional derivative of  $F^k(\theta)$  in direction  $\theta$  is given by (Overton and Womersley, 1993)

$$F^{k'}(\theta; \theta) = \sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} \text{trace}(R^\top L_{ij} R) + \max_{U \in \Phi_{t,k-r}} \sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} \text{trace}(Q^\top L_{ij} Q U), \quad (19)$$

where  $L_{ij} = \partial L(\theta) / \partial \theta_{ij}$ , the matrix  $R \in \mathbb{R}^{N \times r}$  has  $j$ -th column equal to the eigenvector of the  $j$ -th largest eigenvalue of  $L(\theta)$  and the matrix  $Q \in \mathbb{R}^{N \times t}$  has  $j$ -th column equal to the eigenvector of the  $(r+j)$ -th largest eigenvalue of  $L(\theta)$ . In addition the set  $\Phi_{a,b}$  is defined as,

$$\Phi_{a,b} := \{U \in \mathcal{S}_a | U \text{ and } I - U \text{ are positive semi-definite and } \text{trace}(U) = b\}. \quad (20)$$

Overton and Womersley (1993) have shown that  $F^{k'}(\theta; \theta)$  is the sum of the eigenvalues of  $\sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} R^\top L_{ij} R$  plus the sum of the  $k-r$  largest eigenvalues of  $\sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} Q^\top L_{ij} Q$ . Therefore, the directional derivative of  $\lambda_2(L(\theta))$  in the direction  $\theta$  is given by the smallest eigenvalue of  $\sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} Q^\top L_{ij} Q$ , where the matrix  $Q$  is constructed by any complete set of eigenvectors corresponding to the eigenvalue  $\lambda = \lambda_2(L(\theta))$ .



## 4.2 Minimising $\lambda_2(L(\boldsymbol{\theta}))$ and $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$ .

Applying standard gradient descent methods to functions which are almost everywhere differentiable can result in convergence to sub-optimal points (Wolfe, 1972). This occurs when the method for determining the gradient is applied at a point of non-differentiability and results in a direction which is not a descent. In addition, gradients close to points of non-differentiability may be poorly conditioned from a computational perspective leading to poor performance of the optimisation.

The second eigenvalues of the graph Laplacian matrices, while not differentiable everywhere, benefit from the fact that their minimisation tends to lead to a separation from other eigenvalues. Thus a naive gradient descent algorithm tends to perform well. Notice also that if  $u \in \mathbb{R}^N$  with  $\|u\| = 1$  and  $u \perp \mathbf{1}$  is such that  $u^\top L(\boldsymbol{\theta})u = \lambda_2(L(\boldsymbol{\theta}))$  for some  $\boldsymbol{\theta} \in \Theta$ , then for any  $\boldsymbol{\theta}' \in \Theta$  with  $u^\top L(\boldsymbol{\theta}')u < u^\top L(\boldsymbol{\theta})u$  we have  $\lambda_2(L(\boldsymbol{\theta}')) < \lambda_2(L(\boldsymbol{\theta}))$ , since  $u^\top L(\boldsymbol{\theta}')u$  is an upper bound for  $\lambda_2(L(\boldsymbol{\theta}'))$ . Thus even if  $\lambda_2(L(\boldsymbol{\theta}))$  is a repeated eigenvalue, a descent direction for  $u^\top L(\boldsymbol{\theta})u$  is a descent direction for  $\lambda_2(L(\boldsymbol{\theta}))$ , where  $u$  is any corresponding eigenvector. However, this property does not necessarily hold for  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  since the first eigenvector of  $L_{\text{norm}}(\boldsymbol{\theta})$  depends on  $\boldsymbol{\theta}$ , and thus the second eigenvector  $u$  will not necessarily be orthogonal to the first eigenvector of  $L_{\text{norm}}(\boldsymbol{\theta}')$ .

We assume that the similarity function,  $s$ , is Lipschitz continuous and continuously differentiable in  $P$  for each  $i, j$ , and hence the Laplacian matrices  $L(\boldsymbol{\theta})$  and  $L_{\text{norm}}(\boldsymbol{\theta})$  are element-wise Lipschitz continuous and continuously differentiable in  $\boldsymbol{\theta}$ . These conditions are sufficient for the everywhere directional differentiability of  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$ . Our approach for finding locally minimal solutions for  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  can be seen as a simplification of the general method of Overton and Womersley (1993). Our method alternates between a naive application of a standard gradient based optimisation algorithm, in which the simplicity of the second eigenvalue is assumed to hold everywhere along the optimisation path, and a descent step which (in general) decouples the second eigenvalue. We again discuss only  $\lambda_2(L(\boldsymbol{\theta}))$  explicitly, where  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  is analogous. A description of the method is found in Algorithm 1. Notice that upon convergence of a gradient descent algorithm which assumes the simplicity of  $\lambda_2(L(\boldsymbol{\theta}))$ , if  $\lambda_2(L(\boldsymbol{\theta}))$  is simple then the solution is a local minimum, and so the algorithm terminates. If  $\lambda_2(L(\boldsymbol{\theta}))$  is not simple, then the solution may or may not be a local minimum. As we discuss in Section 4.1, if  $\boldsymbol{\theta}$  is such that  $\lambda_2(L(\boldsymbol{\theta}))$  is not simple, then the directional derivative of  $\lambda_2(L(\boldsymbol{\theta}))$  in direction  $\boldsymbol{\theta}$  is given by the smallest eigenvalue of  $\sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} Q^\top L_{ij} Q$ , where  $Q$  is the matrix with columns corresponding to a complete set of eigenvectors for  $\lambda = \lambda_2(L(\boldsymbol{\theta}))$ , and  $L_{ij} = \partial L(\boldsymbol{\theta}) / \partial \theta_{ij}$ . If  $Q^\top L_{ij} Q = \mathbf{0}$  for all  $i = 1, \dots, d-1; j = 1, \dots, l$ , then  $\boldsymbol{\theta}$  is a local minimum and the method terminates, otherwise  $\exists \boldsymbol{\theta} \in \Theta$  s.t.  $\lambda_1 \left( \sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} Q^\top L_{ij} Q \right) < 0$ , and thus  $\boldsymbol{\theta}$  is a descent direction for  $\lambda_2(L(\boldsymbol{\theta}))$ . It is possible to find a locally steepest descent direction by minimising  $\lambda_1 \left( \frac{1}{\|\boldsymbol{\theta}\|} \sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} Q^\top L_{ij} Q \right)$  over  $\boldsymbol{\theta}$ , however the added computational cost associated with this subproblem outweighs the benefit over a simply chosen unit coordinate vector. Notice that the directional derivative of  $\lambda_{k+2}(L(\boldsymbol{\theta}))$  in direction  $\boldsymbol{\theta}$  is given by the  $(k+1)$ -th eigenvalue of  $\sum_{i=1}^{d-1} \sum_{j=1}^l \theta_{ij} Q^\top L_{ij} Q$ , for  $k = 0, 1, \dots, t-1$ , where  $t$  is the multiplicity of the eigenvalue  $\lambda = \lambda_2(L(\boldsymbol{\theta}))$ . Therefore if there exists  $i \in \{1, \dots, d-1\}, j \in \{1, \dots, l\}$  s.t.  $\lambda_t(Q^\top L_{ij} Q) > 0$  and is simple then  $-e_{ij}$ , where  $-e_{ij}$  is the  $(d-1) \times l$  matrix with zeros except in the  $i, j$ -th entry where it takes the value 1, is a descent direction and  $\exists \gamma > 0$  s.t.  $\lambda_2(L(\boldsymbol{\theta} - \gamma' e_{ij})) < \lambda_3(L(\boldsymbol{\theta} - \gamma' e_{ij}))$  for all  $0 < \gamma' < \gamma$ . On the other hand if  $\lambda_1(Q^\top L_{ij} Q) < 0$  and is simple, then  $e_{ij}$  is such a descent direction. If no such pair  $i, j$  exists, then we select  $i, j$  which maximises  $\max\{\lambda_t(Q^\top L_{ij} Q), -\lambda_1(Q^\top L_{ij} Q)\}$  and set  $\boldsymbol{\theta} = -e_{ij}$  if the maximum was determined by the largest eigenvalue and equal to  $e_{ij}$  otherwise.

---

**Algorithm 1:** Minimising  $\lambda_2(L(\boldsymbol{\theta}))$ 

---

1. Initialise  $\boldsymbol{\theta}$ .
  2. Apply gradient based optimisation to  $\lambda_2(L(\boldsymbol{\theta}))$  assuming differentiability
  3. **if**  $\lambda_2(L(\boldsymbol{\theta}))$  is simple **then**  
    **return**  $\boldsymbol{\theta}$
  4. Find  $Q \in \mathbb{R}^{N \times t}$ , a complete set of  $t$  eigenvectors for eigenvalue  $\lambda = \lambda_2(L(\boldsymbol{\theta}))$ .  
    Find  $L_{ij} = \partial L(\boldsymbol{\theta}) / \partial \boldsymbol{\theta}_{ij}$  for  $i = 1, \dots, d-1$ , and  $j = 1, \dots, t$
  5. **if**  $Q^\top L_{ij} Q = \mathbf{0} \ \forall i = 1, \dots, d-1; j = 1, \dots, t$  **then**  
    **return**  $\boldsymbol{\theta}$
  6. **if**  $\exists i \in \{1, \dots, d-1\}; j \in \{1, \dots, t\}$  s.t.  $\lambda_t(Q^\top L_{ij} Q) > 0$  and is simple **then**  
     $\boldsymbol{\theta}' \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}'} \lambda_2(L(\boldsymbol{\theta}'))$  s.t.  $\boldsymbol{\theta}' = \boldsymbol{\theta} - \gamma e_{ij}$ ,  $\gamma > 0$ ,  $\lambda_2(L(\boldsymbol{\theta}'))$  is simple  
    **go to** 2.
  7. **if**  $\exists i \in \{1, \dots, d-1\}; j \in \{1, \dots, t\}$  s.t.  $\lambda_1(Q^\top L_{ij} Q) < 0$  and is simple **then**  
     $\boldsymbol{\theta}' \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}'} \lambda_2(L(\boldsymbol{\theta}'))$  s.t.  $\boldsymbol{\theta}' = \boldsymbol{\theta} + \gamma e_{ij}$ ,  $\gamma > 0$ ,  $\lambda_2(L(\boldsymbol{\theta}'))$  is simple  
    **go to** 2.
  8.  $(I, J) \leftarrow \operatorname{argmax}_{(i,j)} \max\{\lambda_t(Q^\top L_{ij} Q), -\lambda_1(Q^\top L_{ij} Q)\}$
  9. **if**  $\lambda_t(Q^\top L_{IJ} Q) > -\lambda_1(Q^\top L_{IJ} Q)$  **then**  
     $\boldsymbol{\theta}' \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}'} \lambda_2(L(\boldsymbol{\theta}'))$  s.t.  $\boldsymbol{\theta}' = \boldsymbol{\theta} - \gamma e_{IJ}$ ,  $\gamma > 0$   
    **go to** 4.
  10.  $\boldsymbol{\theta}' \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}'} \lambda_2(L(\boldsymbol{\theta}'))$  s.t.  $\boldsymbol{\theta}' = \boldsymbol{\theta} + \gamma e_{IJ}$ ,  $\gamma > 0$   
    **go to** 4.
  11. **end**
- 

### 4.3 Computing Similarities

It is common to define pairwise similarities of points via a decreasing function of the distance between them. That is, for a decreasing function  $k : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ , the similarity function  $s$  may be written,

$$s(P, i, j) = k\left(\frac{d(p_i, p_j)}{\sigma}\right), \quad (21)$$

where  $d(\cdot, \cdot)$  is a metric and  $\sigma > 0$  is a *scaling parameter*. We have found that the projection pursuit method which we propose can be susceptible to outliers when the standard Euclidean distance metric is used, especially in the case of minimising  $\lambda_2(L(\boldsymbol{\theta}))$ . In this subsection we discuss how to embed a balancing constraint into the distance function. By including this balancing mechanism the projection pursuit is steered away from projections which result in only few data being separated from the remainder of the data set.

While the normalisation of the graph cut objective, given in (2), is extremely effective in emphasising balanced partitions in the general spectral clustering problem (von Luxburg, 2007), we have found that in the projection pursuit formulation a further emphasis on balance is sometimes required. This is especially the case in high dimensional applications. Consider the extreme case where  $d > N$ . Then the projection equation,  $V^\top X = P$ , is an underdetermined system of linear equations. Therefore for any desired projected data set  $P$  there exist  $\boldsymbol{\theta} \in \Theta, c \in \mathbb{R} \setminus \{0\}$  s.t.  $V(\boldsymbol{\theta})^\top X = cP$ . In other words

the projected data can be made to have any distribution, up to a scaling constant. In particular we can generally find projections which induce a sufficient separation of a small group of points from the remainder of the data that the normalisation in (2) is inadequate to obtain a balanced partition. We have observed that in practice even for problems of moderate dimension this situation can occur. The importance of including a balancing constraint in the context of projection pursuit for clustering has been observed previously by Zhang et al. (2009) and Pavlidis et al. (2016).

Emphasising balanced partitions is achieved through the use of a compact constraint set  $\Delta$ , which may be defined using the distribution of the projected data set  $P$ . By defining the metric  $d(\cdot, \cdot)$  in such a way that distances between points extending beyond  $\Delta$  are reduced, we increase the similarity of points outside  $\Delta$  with others. If  $P$  is  $l$  dimensional then we define  $\Delta$  as the rectangle  $\Delta = \prod_{i=1}^l \Delta_i$ , where each  $\Delta_i$  is an interval in  $\mathbb{R}$  which is defined using the distribution of the  $i$ -th component of  $P$ . A convenient way of increasing similarities with points lying outside  $\Delta$  is with a transformation  $T_\Delta : \mathbb{R}^l \rightarrow \mathbb{R}^l$ , defined as follows,

$$T_\Delta(y) = (t_{\Delta_1}(y_1), \dots, t_{\Delta_l}(y_l)), \quad (22)$$

$$t_{\Delta_i}(z) := \begin{cases} -\delta \left( \min \Delta_i - z + (\delta(1-\delta))^{\frac{1}{\delta}} \right)^{1-\delta} + \delta (\delta(1-\delta))^{\frac{1-\delta}{\delta}}, & z < \min \Delta_i \\ z - \min \Delta_i, & z \in \Delta_i \\ \delta \left( z - \max \Delta_i + (\delta(1-\delta))^{\frac{1}{\delta}} \right)^{1-\delta} - \delta (\delta(1-\delta))^{\frac{1-\delta}{\delta}} + \text{Diam}(\Delta_i), & z > \max \Delta_i, \end{cases} \quad (23)$$

where  $\delta \in (0, .5]$  is the distance reducing parameter. Each  $t_{\Delta_i}$  is linear on  $\Delta_i$  but has a smaller gradient outside  $\Delta_i$ . As a result we have  $\|T_\Delta(x) - T_\Delta(y)\| \leq \|x - y\|$  for any  $x, y \in \mathbb{R}^l$ , with strict inequality whenever either  $x$  or  $y$  lies outside  $\Delta$ . We define  $T_\Delta$  in this way so that it is continuously differentiable even at the boundaries of  $\Delta$ , and so does not affect the differentiability properties of the similarity function,  $s$ . Figure 1 illustrates how the function  $T_\Delta$  influences distances and similarities in the univariate case.

In the context of projection pursuit it is convenient to define a full dimensional convex constraint set  $\Delta \subset \mathbb{R}^d$  and define the univariate constraint intervals, which we now index by the corresponding projection angles, via the projection of  $\Delta$  onto each  $V(\theta)_i$ . That is,

$$\Delta_{\theta_i} := [\min\{V(\theta)_i^\top x | x \in \Delta\}, \max\{V(\theta)_i^\top x | x \in \Delta\}]. \quad (24)$$

In our implementation, we define  $\Delta$  to be a scaled covariance ellipsoid centered at the mean of the data. The projections of  $\Delta$  are thus given by intervals of the form,

$$\Delta_{\theta_i} = [\mu_{\theta_i} - \beta\sigma_{\theta_i}, \mu_{\theta_i} + \beta\sigma_{\theta_i}], \quad (25)$$

where  $\mu_{\theta_i}$  and  $\sigma_{\theta_i}$  are the mean and standard deviation of the  $i$ -th component of the projected data set  $P(\theta)$  and the parameter  $\beta \geq 0$  determines the width of the projected constraint interval  $\Delta_{\theta_i}$ .


Figure 2 shows two dimensional projections of the 64 dimensional optical recognition of handwritten digits dataset<sup>1</sup>. The leftmost plot shows the PCA projection which is used as initialisation for the projection pursuit. The remaining plots show the projections arising from the minimisation of  $\lambda_2(L(\theta))$  for a variety of values of  $\beta$ . For  $\beta = \infty$ , i.e., an unconstrained projection, the projection pursuit focuses only on a few data and leaves the remainder of the dataset almost unaffected by the projection. Setting  $\beta = 2.5$  causes the projection pursuit to focus on a larger proportion of the tail of the data distribution. Setting  $\beta = 1.5$  however allows the projection pursuit to identify the cluster structure in the data and find a projection which provides a good separation of three of the clusters in the data, i.e., those shown as black, orange and turquoise in the top right plot.

#### 4.4 Correlated and Orthogonal Projections

The formulation of the optimisation problem associated with projection pursuit based on spectral connectivity places no constraints on the projection matrix,  $V$ , except that its columns should have

<sup>1</sup><https://archive.ics.uci.edu/ml/datasets.html>

Figure 1: Effect of  $T_\Delta$  on Distances and Similarities.




(a) The univariate data set  $P$  is plotted against the transformed data  $T_\Delta(P)$ . The point at  $\approx 15$  lies outside  $\Delta$  and its distance to other points, e.g. the point at  $\approx 2$ , is smaller within  $T_\Delta(P)$  (vertical axis) than in  $P$  (horizontal axis). (b) The kernel density estimate of the transformed data  $T_\Delta(P)$  (---) has a stronger bimodal structure, i.e., two well defined clusters, than that of  $P$  (—), which has multiple small modes caused by outliers. The connection between spectral connectivity and density based clustering has been investigated theoretically by Narayanan et al. (2006) showing that the normalised graph cut is asymptotically related to the density on the separating surface. (c) The affinity matrix of the data set  $P$  has a weaker cluster structure than that of  $T_\Delta(P)$ , shown in (d).

unit norm. A common consideration in dimension reduction methods is that the columns in the projection matrix should be orthogonal, i.e.,  $V_i^\top V_j = 0$  for all  $i \neq j$ . In the context of projection pursuit it is common to generate the columns of  $V$  individually, so that orthogonality of the columns can easily be enforced. Huber (1985) proposes first learning  $V_1$  using the original data, and then for each subsequent column the data are first projected into the null space of all the columns learnt so far. An alternative approach (Niu et al., 2011), is to instead project the gradient of the objective into this null space during a gradient based optimisation. Notice, however, that orthogonality in the columns of  $V$  is not essential for the underlying problem. Another common approach (Huber, 1985) is to transform the data after each column is determined in such a way that the columns learned so far are no longer “interesting”, i.e., have low projection index. This does not enforce orthogonality, but rather steers the projection pursuit away from the columns already learned by making them unattractive to the optimisation method.

We propose a different approach which allows us to learn all of the columns of  $V$  simultaneously. This is achieved by introducing an additional term to the objective function which controls the level

Figure 2: Two dimensional projections of optical recognition of handwritten digits dataset arising from the minimisation of  $\lambda_2(L(\boldsymbol{\theta}))$ , for different values of  $\beta$ . In addition, the initialisation through PCA is also shown. The top row of plots shows the true clusters, while the bottom row shows resulting bi-partitions.



of orthogonality, or alternatively correlation, within the projection matrix. In particular, we consider the objective,

$$\min_{\boldsymbol{\theta} \in \Theta} \lambda_2(L(\boldsymbol{\theta})) + \omega \sum_{i \neq j} (V(\boldsymbol{\theta})_i^\top V(\boldsymbol{\theta})_j)^2, \quad (26)$$

or replacing  $\lambda_2(L(\boldsymbol{\theta}))$  with  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  in the normalised case. This approach serves a dual purpose. In the first case, setting  $\omega > 0$  leads to approximately orthogonal projections, without the need to optimise separately over different projection vectors as is standard. Alternatively, setting  $\omega < 0$  leads to approximately perfect correlation, i.e.,  $V(\boldsymbol{\theta})_i \approx \pm V(\boldsymbol{\theta})_j$  for all  $i, j$ . In the latter case the resulting projection is therefore equivalent to a univariate projection. This is similar to simultaneously considering multiple initialisations, and allowing the optimisation procedure to select from them automatically. This is important as the objectives  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  are non-convex, and as a result applying gradient based optimisation can only guarantee convergence to a local optimum.

Notice also that the formulation in Eq. (26) offers computational benefits over the alternative of optimising separately over each projection vector, since the eigenvalues/vectors computed in each function and gradient evaluation need only be computed once for each iteration over the multiple projection dimensions.

## 5 Connection with Maximal Margin Hyperplanes

In this section we establish a connection between the optimal univariate projection for spectral bi-partitioning using the standard Laplacian and large margin separators. In particular, under suitable conditions, as the scaling parameter tends to zero the optimal projection for spectral bi-partitioning converges to the vector admitting the largest margin hyperplane through the data. This establishes a theoretical connection between spectral connectedness and separability of the resulting clusters in terms of Euclidean distance. Large margin separators are ubiquitous in the machine learning literature, and were first introduced in the context of supervised classification via support vector machines (SVM, Vapnik and Kotz (1982)). In more recent years they have shown to be very useful for unsupervised partitioning in the context of maximum margin clustering as well (Xu et al., 2004;

Zhang et al., 2009).

Our result pertains to univariate projections, and therefore the  $d \times 1$  projection matrix is equivalently viewed as a *projection vector* in  $\mathbb{R}^d$ . We therefore use the notation  $v(\theta)$ , instead of  $V(\theta)$  as before.

The result holds for all similarities for which the function  $k$ , in Eq. (21), satisfies the tail condition  $\lim_{x \rightarrow \infty} k(x + \epsilon)/k(x) = 0$  for all  $\epsilon > 0$ . This condition is satisfied by functions with exponentially decaying tails, including the popular Gaussian and Laplace kernels. It is, however, not satisfied by those with polynomially decaying tails.

The constraint set  $\Delta$  again plays an important role as in many cases the largest margin hyperplane through a set of data separates only a few points from the rest, making it meaningless for the purpose of clustering. We therefore prefer to restrict the hyperplane to intersect the set  $\Delta$ . What we in fact show in this section is that there exists a set  $\Delta' \subset \Delta$  satisfying  $\Delta' \cap X = \Delta \cap X$ , such that, as the scaling parameter tends to zero, the optimal projection for  $\lambda_2(L(\theta))$  converges to the projection admitting the largest margin hyperplane that intersects  $\Delta'$ . The distinction between the largest margin hyperplane intersecting  $\Delta'$  and that intersecting  $\Delta$  is scarcely of practical relevance, but plays an important role in the theory we present in this section. It accounts for situations when the largest margin hyperplane intersecting  $\Delta$  lies close to its boundary and the distance between the hyperplane and the nearest point outside  $\Delta$  is larger than to the nearest point inside  $\Delta$ . Aside from this very specific case, the two in fact coincide.

A hyperplane is a translated subspace of co-dimension 1, and can be parameterised by a non-zero vector  $v \in \mathbb{R}^d \setminus \{0\}$  and scalar  $b$  as the set  $H(v, b) = \{x \in \mathbb{R}^d \mid v^\top x = b\}$ . Clearly, for any  $c \in \mathbb{R} \setminus \{0\}$ , one has  $H(v, b) = H(cv, cb)$ , and so we can assume that  $v$  has unit norm, thus the same parameterisation by  $\theta$  can be used. For a finite set of points  $X \subset \mathbb{R}^d$ , the *margin* of hyperplane  $H(v(\theta), b)$  w.r.t.  $X$  is the minimal Euclidean distance between  $H(v(\theta), b)$  and  $X$ . That is,

$$\text{margin}(v(\theta), b) = \min_{x \in X} |v(\theta)^\top x - b|. \quad (27)$$

Connections between maximal margin hyperplanes and Bayes optimal hyperplanes as well as minimum density hyperplanes have been established (Tong and Koller, 2000; Pavlidis et al., 2016).

In this section we will use the notation  $v^\top X = \{v^\top x_1, \dots, v^\top x_N\}$ , and for a set  $P \subset \mathbb{R}$  and  $y \in \mathbb{R}$  we write, for example,  $P_{>y}$  for  $P \cap (y, \infty)$ . For scaling parameter  $\sigma > 0$  and distance reduction factor  $\delta > 0$  define  $\theta_{\sigma, \delta} := \arg\min_{\theta \in \Theta} \lambda_2(L(\theta, \sigma, \delta))$ , where  $L(\theta, \sigma, \delta)$  is as  $L(\theta)$  from before, but with an explicit dependence on the scaling parameter and distance reducing parameter used in the similarity function. That is,  $\theta_{\sigma, \delta}$  defines the projection generating the minimal spectral connectivity of  $X$  for a given pair  $\sigma, \delta$ .

Before proving the main result of this section, we require the following supporting results. Lemma 2 provides a lower bound on the second eigenvalue of the graph Laplacian of a one dimensional data set in terms of the largest Euclidean separation of adjacent points, with respect to a constraint set  $\Delta$ . This lemma also shows how we construct the set  $\Delta'$ . Lemma 3 uses this result to show that a projection angle  $\theta \in \Theta$  leads to lower spectral connectivity than all projections admitting smaller maximal margin hyperplanes intersecting  $\Delta'$  for all pairs  $\sigma, \delta$  sufficiently close to zero.

**Lemma 2** *Let  $k : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  be a non-increasing, positive function and let  $\sigma > 0, \delta \in (0, 0.5]$ . Let  $P = \{p_1, \dots, p_N\}$  be a univariate data set and let  $\Delta = [a, b]$  for  $a < b \in \mathbb{R}$ . Suppose that  $|P \cap \Delta| \geq 2$  and  $a \geq \min\{P\}, b \leq \max\{P\}$ . Define  $\Delta' = [a', b']$ , where  $a' = (a + \min\{P \cap \Delta\})/2, b' = (b + \max\{P \cap \Delta\})/2$ . Let  $M = \max_{x \in \Delta'} \{\min_{i=1 \dots N} |x - p_i|\}$ . Define  $L(P)$  to be the Laplacian of the graph with vertices  $P$  and similarities according to  $s(P, i, j) = k(|T_\Delta(p_i) - T_\Delta(p_j)|/\sigma)$ . Then  $\lambda_2(L(P)) \geq \frac{1}{|P|^3} k((2M + \delta C)/\sigma)$ , where  $C = \max\{D, D^{1-\delta}\}, D = \max\{a - \min\{P\}, \max\{P\} - b\}$ .*

**Proof** We can assume that  $P$  is sorted in increasing order, i.e.  $p_i \leq p_{i+1}$ , since this does not affect

the eigenvalues of  $L(P)$ . We first show that  $s(P, i, i+1) \geq k((2M + \delta C)/\sigma)$  for all  $i = 1, \dots, N-1$ . To this end observe that  $\delta \left( x + \left( \delta(1 - \delta)^{\frac{1}{\delta}} \right)^{1-\delta} - \delta(\delta(1 - \delta))^{\frac{1-\delta}{\delta}} \right) \leq \delta \max\{x, x^{1-\delta}\}$  for  $x \geq 0$ .

- If  $p_i, p_{i+1} \leq a$  then  $s(P, i, i+1) = k((T_\Delta(p_{i+1}) - T_\Delta(p_i))/\sigma) \geq k((T_\Delta(a) - T_\Delta(p_i))/\sigma) \geq k((2M + \delta C)/\sigma)$  by the definition of  $C$  and using the above inequality, since  $k$  is non-increasing. The case  $p_i, p_{i+1} \geq b$  is similar.
- If  $p_i, p_{i+1} \in \Delta$  then  $p_i, p_{i+1} \in \Delta' \Rightarrow |p_i - p_{i+1}| \leq 2M \Rightarrow s(P, i, i+1) \geq k(2M/\sigma) \geq k((2M + \delta C)/\sigma)$  since  $M$  is the largest margin in  $\Delta'$ .
- If none the above hold, then we lose no generality in assuming  $p_i < a$ ,  $a < p_{i+1} < b$  since the case  $a < p_i < b$ ,  $p_{i+1} > b$  is analogous. We must have  $p_{i+1} = \min\{P \cap \Delta\}$  and so  $a' = (a + p_{i+1})/2$ . If  $p_{i+1} - a > 2M$  then  $\min_{j=1 \dots N} |a' - p_j| > M$ , a contradiction since  $a' \in \Delta'$  and  $M$  is the largest margin in  $\Delta'$ . Therefore  $p_{i+1} - a \leq 2M$ . In all  $T_\Delta(p_{i+1}) - T_\Delta(p_i) = (p_{i+1} - a) + \delta(a - p_i + (\delta(1 - \delta))^{\frac{1}{\delta}})^{1-\delta} - \delta(\delta(1 - \delta))^{\frac{1-\delta}{\delta}} \leq 2M + \delta C \Rightarrow s(P, i, i+1) \geq k((2M + \delta C)/\sigma)$ .

Now, let  $u$  be the second eigenvector of  $L(P)$ . Then  $\|u\| = 1$  and  $u \perp \mathbf{1}$  and therefore  $\exists i, j$  s.t.  $u_i - u_j \geq \frac{1}{\sqrt{|P|}}$ . We thus know that there exists  $m$  s.t.  $|u_m - u_{m+1}| \geq \frac{1}{|P|^{3/2}}$ . By (von Luxburg, 2007, Proposition 1), we know that  $u^\top L(P)u = \frac{1}{2} \sum_{i,j} s(P, i, j)(u_i - u_j)^2 \geq s(P, m, m+1)(u_m - u_{m+1})^2 \geq \frac{1}{|P|^3} k((2M + \delta C)/\sigma)$  since all consecutive pairs  $p_m, p_{m+1}$  have similarity at least  $k((2M + \delta C)/\sigma)$ , by above. Therefore  $\lambda_2(L(P)) \geq \frac{1}{|P|^3} k((2M + \delta C)/\sigma)$  as required.

In the above Lemma we have assumed that  $\Delta$  is contained within the convex hull of the points  $P$ , however the results of this section can easily be modified to allow for cases where this does not hold. In particular, if an unconstrained large margin hyperplane is sought, then setting  $\Delta$  to be arbitrarily large allows for this. We have merely stated the results in the most convenient context for our practical implementation.

The set  $\Delta'$  in the above is defined in terms of the one dimensional constraint set  $[a, b]$ . We define the full dimensional set  $\Delta'$  along the same lines by,

$$\begin{aligned} \Delta' &:= \{x \in \mathbb{R}^d | v(\theta)^\top x \in \Delta'_\theta \ \forall \theta \in \Theta\}, \\ \Delta'_\theta &= \left[ \frac{\min \Delta_\theta + \min\{v(\theta)^\top X \cap \Delta_\theta\}}{2}, \frac{\max \Delta_\theta + \max\{v(\theta)^\top X \cap \Delta_\theta\}}{2} \right]. \end{aligned} \quad (28)$$

Here we assume that  $\Delta$  is contained within the convex hull of the  $d$ -dimensional data set  $X$ . Notice that since  $\Delta$  is convex, we have  $v(\theta)^\top \Delta' = \Delta'_\theta$ . In what follows we show that as  $\sigma$  and  $\delta$  are reduced to zero the optimal projection for spectral partitioning converges to the projection admitting the largest margin hyperplane intersecting  $\Delta'$ . If it is the case that the largest margin hyperplane intersecting  $\Delta$  also intersects  $\Delta'$ , as is often the case, although this fact will not be known, then it is actually not necessary that  $\delta$  tend towards zero. In such cases it only needs to satisfy  $\delta \leq 2M/C$  for the corresponding values of  $M$  and  $C$  over all possible projections. In particular, choosing  $\max\{\text{Diam}(X), \text{Diam}(X)^{1-\delta}\}$  instead of  $C$  is appropriate for all projections.

**Lemma 3** Let  $\theta \in \Theta$  and let  $k: \mathbb{R}_+ \rightarrow \mathbb{R}_+$  be non-increasing, positive, and satisfy

$$\lim_{x \rightarrow \infty} k(x(1 + \epsilon))/k(x) = 0$$

for all  $\epsilon > 0$ . Then for any  $0 < m < \max_{b \in \Delta'_\theta} \text{margin}(v(\theta), b)$  there exists  $\sigma' > 0$  and  $\delta' > 0$  s.t.  $0 < \sigma < \sigma'$ ,  $0 < \delta < \delta'$  and  $\max_{c \in \Delta'_\theta} \text{margin}(v(\theta'), c) < \max_{b \in \Delta'_\theta} \text{margin}(v(\theta), b) - m \Rightarrow \lambda_2(L(\theta, \sigma, \delta)) < \lambda_2(L(\theta', \sigma, \delta))$ .

**Proof** Let  $B = \operatorname{argmax}_{b \in \Delta'_\theta} \operatorname{margin}(v(\theta), b)$  and  $M = \operatorname{margin}(v(\theta), B)$ . We assume that  $M \neq 0$ , since otherwise there is nothing to show. Now, since spectral clustering solves a relaxation of the minimum normalised cut problem we have,

$$\begin{aligned}
\lambda_2(L(\theta, \sigma, \delta)) &\leq \frac{1}{|X|} \min_{C \subset X} \sum_{\substack{i,j: x_i \in C \\ x_j \notin C}} s(P(\theta), i, j) \left( \frac{1}{|C|} + \frac{1}{|X \setminus C|} \right) \\
&\leq \frac{1}{|X|} \sum_{\substack{i,j: v(\theta)^\top x_i < B \\ v(\theta)^\top x_j > B}} s(P(\theta), i, j) \left( \frac{1}{|(v(\theta)^\top X)_{<B}|} + \frac{1}{|(v(\theta)^\top X)_{>B}|} \right) \\
&= \frac{1}{|X|} \sum_{\substack{i,j: v(\theta)^\top x_i < B \\ v(\theta)^\top x_j > B}} k \left( \frac{T_{\Delta_\theta}(v(\theta)^\top x_j) - T_{\Delta_\theta}(v(\theta)^\top x_i)}{\sigma} \right) \left( \frac{|X|}{|(v(\theta)^\top X)_{<B}| |(v(\theta)^\top X)_{>B}|} \right) \\
&\leq |(v(\theta)^\top X)_{<B}| |(v(\theta)^\top X)_{>B}| k \left( \frac{2M}{\sigma} \right) \left( \frac{1}{|(v(\theta)^\top X)_{<B}| |(v(\theta)^\top X)_{>B}|} \right) \\
&= k(2M/\sigma).
\end{aligned}$$

The final inequality holds since for any  $i, j$  s.t.  $v(\theta)^\top x_i < B$  and  $v(\theta)^\top x_j > B$  we must have  $T_{\Delta_\theta}(v(\theta)^\top x_j) - T_{\Delta_\theta}(v(\theta)^\top x_i) \geq 2M$ . Now, for any  $\theta' \in \Theta$ , let  $M_{\theta'} = \max_{c \in \Delta'_{\theta'}} \operatorname{margin}(v(\theta'), c)$ . By Lemma 2 we know that  $\lambda_2(L(\theta', \sigma, \delta)) \geq \frac{1}{|X|^3} k((2M_{\theta'} + \delta C)/\sigma)$ , where  $C = \max\{\operatorname{Diam}(X), \operatorname{Diam}(X)^{1-\delta}\}$ . Therefore,

$$\begin{aligned}
\lim_{\sigma \rightarrow 0^+, \delta \rightarrow 0^+} \frac{\lambda_2(L(\theta, \sigma, \delta))}{\inf_{\theta' \in \Theta} \{\lambda_2(L(\theta', \sigma, \delta)) \mid M_{\theta'} < M - m\}} &\leq \lim_{\sigma \rightarrow 0^+, \delta \rightarrow 0^+} \frac{|X|^3 k(2M/\sigma)}{k((2(M - m) + \delta C)/\sigma)} \\
&= 0.
\end{aligned}$$

This gives the result.

Lemma 3 shows almost immediately that the margin admitted by the optimal projection for spectral bi-partitioning converges to the largest margin through  $\Delta'$  as  $\sigma$  and  $\delta$  go to zero. The main result of this section, Theorem 4, shows the stronger result that the optimal projection itself converges to the projection admitting the largest margin.

**Theorem 4** Let  $X = \{x_1, \dots, x_N\}$  and suppose that there is a unique hyperplane, which can be parameterised by  $(v(\theta^*), b^*)$ , intersecting  $\Delta'$  and attaining maximal margin on  $X$ . Let  $k: \mathbb{R}_+ \rightarrow \mathbb{R}_+$  be decreasing, positive and satisfy  $\lim_{x \rightarrow \infty} k((1 + \epsilon)x)/k(x) = 0$  for all  $\epsilon > 0$ . Then,

$$\lim_{\sigma \rightarrow 0^+, \delta \rightarrow 0^+} v(\theta_{\sigma, \delta}) = v(\theta^*).$$


**Proof** Take any  $\epsilon > 0$ . Pavlidis et al. (2016) have shown that  $\exists m_\epsilon > 0$  s.t. for  $w \in \mathbb{R}^d, c \in \mathbb{R}$ ,  $\|(w, c)/\|w\| - (v(\theta^*), b^*)\| > \epsilon \Rightarrow \operatorname{margin}(w/\|w\|, c/\|w\|) < \operatorname{margin}(v(\theta^*), b^*) - m_\epsilon$ . By Lemma 3 we know  $\exists \sigma' > 0, \delta' > 0$  s.t. if  $0 < \sigma < \sigma'$  and  $0 < \delta < \delta'$  then  $\exists c \in \Delta_\theta$  s.t.  $\operatorname{margin}(v(\theta_{\sigma, \delta}), c) \geq \operatorname{margin}(v(\theta^*), b^*) - m_\epsilon$ , since  $\theta_{\sigma, \delta}$  is optimal for the pair  $\sigma, \delta$ . Thus, by above,  $\|(v(\theta_{\sigma, \delta}), c) - (v(\theta^*), b^*)\| \leq \epsilon$ . But  $\|(v(\theta_{\sigma, \delta}), c) - (v(\theta^*), b^*)\| \geq \|v(\theta_{\sigma, \delta}) - v(\theta^*)\|$  for any  $c \in \mathbb{R}$ . Since  $\epsilon > 0$  was arbitrary, we therefore have  $v(\theta_{\sigma, \delta}) \rightarrow v(\theta^*)$  as  $\sigma, \delta \rightarrow 0^+$ .

While the results of this section apply only for univariate projections, we have observed empirically that if a shrinking sequence of scaling parameters is employed for a multivariate projection, then the projected data tend to display large Euclidean separation. This is illustrated in Figure 3 which shows two dimensional plots of the 72 dimensional yeast cell cycle analysis dataset<sup>2</sup>. As in Figure 2 the

<sup>2</sup><http://genome-www.stanford.edu/cellcycle/>



Figure 3: Large Euclidean separation of the yeast cell cycle dataset. The left plots show the result from a 2 dimensional projection pursuit using the proposed method. The middle plots show the 1 dimensional projection pursuit result. The right plots show the result of the maximum margin clustering method of Zhang et al. (2009).



top plots show the true clusters in the data and the bottom plots show the clustering assignments. The left plots show the result of a two dimensional projection pursuit using the proposed method. In the middle plots the first projection is learnt using one dimensional projection pursuit, and the second is set to be the direction of maximum variance within the null space of the first projection. The right plots use as the first projection the result of the iterative support vector regression method for maximum margin clustering (Zhang et al., 2009), and again the second projection is the direction of maximum variance within its null space.

A similar intuition which underlies the theoretical results of this section can be used to reason why this will occur in multivariate projections, in that as the scaling parameter reduces to zero the value of  $\lambda_2(L(\boldsymbol{\theta}))$  is controlled by the smallest distance between observations in different elements of the induced partition. It is however elusive how to formulate this rigorously in the presence of the constraint set  $\Delta$  in more than one dimension.

## 6 Speeding up Computation using Microclusters

In this section we discuss how a preprocessing of the data using *microclusters* can be used to significantly speed up the optimisation process. We derive theoretical bounds on the error induced by this approximation. Our approach uses a result from matrix perturbation theory for diagonally dominant matrices, and therefore only applies to the standard Laplacian,  $L(\boldsymbol{\theta})$ . However, we have seen empirically that a close approximation of the optimisation surface is obtained for both  $\lambda_2(L(\boldsymbol{\theta}))$  and  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$ .

The concept of a microcluster was introduced by Zhang et al. (1996) in the context of clustering very large data sets. Microclusters are small clusters of data which can in turn be clustered to generate a clustering of the entire data set. A microcluster like approach in the context of spectral clustering has been considered by Yan et al. (2009), where the authors obtain bounds on the mis-clustering rate induced by the approximation. Rather than using microclusters as an intermediate step towards determining a final clustering model, we use them to form an approximation of the optimisation surface for projection pursuit which is less computationally expensive to explore. The error bound

depends on the ratio of cluster radii to scaling parameter. As such, this method does not provide a good approximation when  $\sigma$  is close to zero. Our bounds rely on the following result from perturbation theory.

**Theorem 5** *Ye (2009)*

Let  $A = [a_{ij}]$  and  $\tilde{A} = [\tilde{a}_{ij}]$  be two symmetric positive semidefinite diagonally dominant matrices, and let  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$  and  $\tilde{\lambda}_1 \leq \tilde{\lambda}_2 \leq \dots \leq \tilde{\lambda}_n$  be their respective eigenvalues. If, for some  $0 \leq \epsilon < 1$ ,  $|a_{ij} - \tilde{a}_{ij}| \leq \epsilon |a_{ij}| \forall i \neq j$ , and  $|v_i - \tilde{v}_i| \leq \epsilon v_i \forall i$ , where  $v_i = a_{ii} - \sum_{j \neq i} |a_{ij}|$ , and similarly for  $\tilde{v}_i$ , then

$$|\lambda_i - \tilde{\lambda}_i| \leq \epsilon \lambda_i \forall i.$$

An inspection of the proof of Theorem 5 reveals that  $\epsilon < 1$  is necessary only to ensure that the signs of  $a_{ij}$  are the same as those of  $\tilde{a}_{ij}$ . In the case of Laplacian matrices this equivalence of signs holds by design, and so in this context the requirement that  $\epsilon < 1$  can be relaxed.

In the microcluster approach, the data set  $X = \{x_1, \dots, x_N\}$  is replaced with  $K$  points  $c_1, \dots, c_K$  which represent the centers of a  $K$ -clustering of  $X$ . By projecting these microcluster centers during subspace optimisation, rather than the data themselves, the computational cost associated with each eigen problem is reduced from  $\mathcal{O}(N^2)$  to  $\mathcal{O}(K^2)$ . If we define the radius,  $\rho$ , of a cluster  $C$  to be the greatest distance between one of its members and its center, that is,

$$\rho(C) = \max_{x \in C} \left\| x - \frac{1}{|C|} \sum_{x \in C} x \right\|, \quad (29)$$

then we expect the approximation error to be small whenever the microcluster radii are small. The bounds on the approximation error which we present in this section are worst case and rely on standard eigenvalue bounds, and so can be pessimistic. To obtain a reasonable bound on the approximation surface, as many as  $K \approx 0.6N$  might be needed, leading to only a threefold speed up. We have observed empirically, however, that even for  $K = 0.1N$  (and sometimes lower) one still obtains a close approximation of the optimisation surface. This makes the projection pursuit of the order of 100 times faster.

**Lemma 6** *Let  $\mathcal{C} = C_1, \dots, C_K$  be a  $K$ -clustering of  $X$  with centers  $c_1, \dots, c_K$ , radii  $\rho_1, \dots, \rho_K$  and counts  $n_1, \dots, n_K$ . For  $\boldsymbol{\theta} \in \Theta$  define  $N(\boldsymbol{\theta}), B(\boldsymbol{\theta}) \in \mathbb{R}^{K \times K}$  where  $N(\boldsymbol{\theta})$  is the diagonal matrix with*

$$N(\boldsymbol{\theta})_{i,i} = \sum_{j=1}^K n_j s(P^c(\boldsymbol{\theta}), i, j)$$

and

$$B(\boldsymbol{\theta})_{i,j} = \sqrt{n_i n_j} s(P^c(\boldsymbol{\theta}), i, j),$$

where  $P^c(\boldsymbol{\theta}) = \{V(\boldsymbol{\theta})^\top c_1, \dots, V(\boldsymbol{\theta})^\top c_K\}$  are the projected microcluster centers and the similarities are given by  $s(P^c(\boldsymbol{\theta}), i, j) = k(\|T_{\Delta_{\boldsymbol{\theta}}}(V(\boldsymbol{\theta})^\top c_i) - T_{\Delta_{\boldsymbol{\theta}}}(V(\boldsymbol{\theta})^\top c_j)\|/\sigma)$ , and  $k(x)$  is positive and non-increasing for  $x \geq 0$ . Then,

$$\frac{|\lambda_2(L(\boldsymbol{\theta})) - \lambda_2(N(\boldsymbol{\theta}) - B(\boldsymbol{\theta}))|}{\lambda_2(L(\boldsymbol{\theta}))} \leq \max_{i \neq j} \max \left\{ 1 - \frac{k(D_{ij}/\sigma)}{k((D_{ij} - \rho_i - \rho_j)^+/\sigma)}, \frac{k(D_{ij}/\sigma)}{k((D_{ij} + \rho_i + \rho_j)/\sigma)} - 1 \right\},$$

where  $D_{ij} = \|T_{\Delta_{\boldsymbol{\theta}}}(V(\boldsymbol{\theta})^\top c_i) - T_{\Delta_{\boldsymbol{\theta}}}(V(\boldsymbol{\theta})^\top c_j)\|$  and  $(x)^+ = \max\{0, x\}$ .

**Proof** For brevity we temporarily drop the notational dependence on  $\theta$ . Let  $P^{c'} = \{V^\top c_1, V^\top c_1, \dots, V^\top c_K, V^\top c_K\}$ , where each  $V^\top c_i$  is repeated  $n_i$  times. Let  $L^{c'}$  be the Laplacian of the graph with vertices  $P^{c'}$  and edges given by  $s(P^{c'}, i, j)$ . We begin by showing that  $\lambda_2(L^{c'}) = \lambda_2(N - B)$ . Take  $v \in \mathbb{R}^K$ , then,

$$\begin{aligned} v^\top (N - B)v &= \sum_{i,j} s(P^{c'}, i, j)(v_i^2 n_j - v_i v_j \sqrt{n_i n_j}) \\ &= \frac{1}{2} \sum_{i,j} s(P^{c'}, i, j)(v_i^2 n_j + v_j^2 n_i - 2v_i v_j \sqrt{n_i n_j}) \\ &\geq 0, \end{aligned}$$

and so  $N - B$  is positive semi-definite. In addition, it is straightforward to verify that  $(N - B)(\sqrt{n_1} \dots \sqrt{n_K}) = \mathbf{0}$ , and hence 0 is the smallest eigenvalue of  $N - B$  with eigenvector  $(\sqrt{n_1} \dots \sqrt{n_K})$ . Now, let  $u$  be the second eigenvector of  $L^{c'}$ . Then  $u_j = u_k$  for pairs of indices  $j, k$  aligned with the same  $V^\top c_i$  in  $P^{c'}$ . Define  $u^c \in \mathbb{R}^K$  s.t.  $u_i^c = \sqrt{n_i} u_j$  where index  $j$  is aligned with  $V^\top c_i$  in  $P^{c'}$ . Then  $(u^c)^\top (\sqrt{n_1} \dots \sqrt{n_K}) = \sum_{i=1}^K u_i^c \sqrt{n_i} = \sum_{i=1}^K n_i u_{j_i}$  where index  $j_i$  is aligned with  $V^\top c_i$  in  $P^{c'}$  for each  $i$ . Therefore  $n_i u_{j_i} = \sum_{j: P^{c'} = V^\top c_i} u_j$  and hence  $(u^c)^\top (\sqrt{n_1} \dots \sqrt{n_K}) = \sum_{i=1}^K \sum_{j: P^{c'} = V^\top c_i} u_j = \sum_{i=1}^N u_i = 0$  since  $\mathbf{1}$  is the smallest eigenvector of  $L^{c'}$  and so  $u \perp \mathbf{1}$ . Similarly  $\|u^c\|^2 = \sum_{i=1}^K n_i u_{j_i}^2 = \sum_{i=1}^N u_i^2 = 1$ . Thus  $u^c \perp (\sqrt{n_1} \dots \sqrt{n_K})$  and  $\|u^c\| = 1$  and so is a candidate for the second eigenvector of  $N - B$ . In addition it is straightforward to show that  $(u^c)^\top (N - B)u^c = u \cdot L^{c'}u$ . Now, suppose by way of contradiction that  $\exists w \perp (\sqrt{n_1} \dots \sqrt{n_K})$  with  $\|w\| = 1$  s.t.  $w^\top (N - B)w < (u^c)^\top (N - B)u^c$ . Then let  $w' = (w_1/\sqrt{n_1} \dots w_K/\sqrt{n_K})$  where each  $w_i/\sqrt{n_i}$  is repeated  $n_i$  times. Then  $\|w'\| = 1$ ,  $(w')^\top \mathbf{1} = w^\top (\sqrt{n_1} \dots \sqrt{n_K}) = 0$  and  $w^\top L^{c'}w < u^\top L^{c'}u$ , a contradiction since  $u$  is the second eigenvector of  $L^{c'}$ .

Now, let  $i, j, m, n$  be such that  $x_m \in C_i$  and  $x_n \in C_j$ . We temporarily drop the notational dependence on  $\Delta$ . Then,

$$\begin{aligned} \|T(V^\top x_m) - T(V^\top x_n)\| &= \|T(V^\top x_m) - T(V^\top c_i) + T(V^\top c_i) - T(V^\top c_j) \\ &\quad + T(V^\top c_j) - T(V^\top x_n)\| \\ &\leq \|T(V^\top x_m) - T(V^\top c_i)\| + \|T(V^\top c_i) - T(V^\top c_j)\| \\ &\quad + \|T(V^\top c_j) - T(V^\top x_n)\| \\ &\leq \rho_i + \rho_j + D_{ij}, \end{aligned}$$

since  $T$  contracts distances and  $\rho_i$  and  $\rho_j$  are the radii of  $C_i$  and  $C_j$ . Since  $k$  is non-increasing we therefore have,

$$\begin{aligned} \frac{k(D_{ij}/\sigma)}{k((D_{ij} - \rho_i - \rho_j)^+/\sigma)} &\leq \frac{k(D_{ij}/\sigma)}{k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)} \leq \frac{k(D_{ij}/\sigma)}{k((D_{ij} + \rho_i + \rho_j)/\sigma)} \\ \Rightarrow 1 - \frac{k(D_{ij}/\sigma)}{k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)} &\leq 1 - \frac{k(D_{ij}/\sigma)}{k((D_{ij} - \rho_i - \rho_j)^+/\sigma)} \end{aligned}$$

and

$$\frac{k(D_{ij}/\sigma)}{k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)} - 1 \leq \frac{k(D_{ij}/\sigma)}{k((D_{ij} + \rho_i + \rho_j)/\sigma)} - 1.$$

Therefore

$$\left| \frac{k(D_{ij}/\sigma)}{k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)} - 1 \right| \leq \max \left\{ 1 - \frac{k(D_{ij}/\sigma)}{k((D_{ij} - \rho_i - \rho_j)^+/\sigma)}, \frac{k(D_{ij}/\sigma)}{k((D_{ij} + \rho_i + \rho_j)/\sigma)} - 1 \right\}.$$

Now, we lose no generality by assume that  $X$  is ordered such that for each  $i$  the elements of cluster  $C_i$  are aligned with  $V^\top c_i$  in  $P^{c'}$ , since this does not affect the eigenvalues of the Laplacian of  $V^\top X$ ,  $L$ . By the design of the Laplacian matrix the “ $v_i$ ” of Theorem 5 are exactly zero. For off diagonal terms  $m, n$  with corresponding  $i, j$  as above, consider

$$\begin{aligned} \frac{|L_{mn} - L'_{mn}|}{|L_{mn}|} &= \frac{|k(D_{ij}/\sigma) - k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)|}{k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)} \\ &= \left| \frac{k(D_{ij}/\sigma)}{k(\|T(V^\top x_m) - T(V^\top x_n)\|/\sigma)} - 1 \right|. \end{aligned}$$

Theorem 5 thus gives the result.

The above bound depends on  $\theta$  via the quantity  $D_{ij}$  and for some functions  $k$  it is difficult to remove this dependence. We consider the class of functions, parameterised by  $\alpha \geq 0$ , and given by

$$k(x) = \left( \frac{|x|}{\alpha} + 1 \right)^\alpha \exp(-|x|), \quad (30)$$

where we adopt the convention  $(\frac{a}{0})^0 = 1$  for any  $a \in \mathbb{R}$ . For  $\alpha = 0$  this is equivalent to the Laplace kernel, but for  $\alpha > 0$  has the useful property of being differentiable at 0. We have found the choice of  $k$  to matter little in the results of the proposed approach. The above class of functions is chosen as it allows us to obtain a uniform bound on the error induced by the above approximation. Note the parameter  $\alpha$  is not intended as a tuning parameter, but rather we set  $\alpha$  close to zero to obtain a function similar to the Laplace kernel, but which is differentiable at zero.

**Corollary 7** *Let the conditions of Lemma 6 hold, and let  $k(x) = (|x|/\alpha + 1)^\alpha \exp(-|x|)$  for  $\alpha \geq 0$ . Then,*

$$\frac{|\lambda_2(L(\theta)) - \lambda_2(N(\theta) - B(\theta))|}{\lambda_2(L(\theta))} \leq \max_{i \neq j} \left( \frac{\text{Diam}(X) + \sigma\alpha}{\text{Diam}(X) + \rho_i + \rho_j + \sigma\alpha} \right)^\alpha \exp\left(\frac{\rho_i + \rho_j}{\sigma}\right) - 1.$$

**Proof** Firstly, consider

$$\frac{k(D_{ij}/\sigma)}{k((D_{ij} + \rho_i + \rho_j)/\sigma)} - 1 = \left( \frac{D_{ij} + \sigma\alpha}{D_{ij} + \rho_i + \rho_j + \sigma\alpha} \right)^\alpha \exp\left(\frac{\rho_i + \rho_j}{\sigma}\right) - 1.$$

Now, the function  $\left( \frac{x + \sigma\alpha}{x + y + \sigma\alpha} \right)^\alpha \exp(y/\sigma)$  is non-decreasing in  $x$  for  $x, y, \alpha, \sigma \geq 0$ , therefore by above

$$\frac{k(D_{ij}/\sigma)}{k((D_{ij} + \rho_i + \rho_j)/\sigma)} - 1 \leq \left( \frac{\text{Diam}(X) + \sigma\alpha}{\text{Diam}(X) + \rho_i + \rho_j + \sigma\alpha} \right)^\alpha \exp\left(\frac{\rho_i + \rho_j}{\sigma}\right) - 1.$$

Secondly, consider the case  $D_{ij} \geq \rho_i + \rho_j$ , then

$$\begin{aligned} 1 - \frac{k(D_{ij}/\sigma)}{k((D_{ij} - \rho_i - \rho_j)^+/\sigma)} &= 1 - \left( \frac{D_{ij} + \sigma\alpha}{D_{ij} - \rho_i - \rho_j + \sigma\alpha} \right)^\alpha \exp\left(-\frac{\rho_i + \rho_j}{\sigma}\right) \\ &\leq 1 - \left( \frac{\text{Diam}(X) + \rho_i + \rho_j + \sigma\alpha}{\text{Diam}(X) + \sigma\alpha} \right)^\alpha \exp\left(-\frac{\rho_i + \rho_j}{\sigma}\right), \end{aligned}$$

since  $\left( \frac{x + \sigma\alpha}{x - y + \sigma\alpha} \right)^\alpha \exp(y/\sigma)$  is non-increasing in  $x$  for  $x, y, \alpha, \sigma \geq 0$ . On the other hand, if  $D_{ij} < \rho_i + \rho_j$  then,

$$\begin{aligned} 1 - \frac{k(D_{ij}/\sigma)}{k((D_{ij} - \rho_i - \rho_j)^+/\sigma)} &= 1 - k\left(\frac{D_{ij}}{\sigma}\right) \leq 1 - k\left(\frac{\rho_i + \rho_j}{\sigma}\right) \\ &= 1 - \left( \frac{\rho_i + \rho_j + \sigma\alpha}{\sigma\alpha} \right)^\alpha \exp\left(-\frac{\rho_i + \rho_j}{\sigma}\right) \\ &\leq 1 - \left( \frac{\text{Diam}(X) + \rho_i + \rho_j + \sigma\alpha}{\text{Diam}(X) + \sigma\alpha} \right)^\alpha \exp\left(-\frac{\rho_i + \rho_j}{\sigma}\right), \end{aligned}$$

where the first inequality comes from the fact that  $D_{ij} < \rho_i + \rho_j$  and  $k$  is decreasing. Now, using the identity  $1 - \frac{1}{x} \leq x - 1$  for  $x \neq 0$ , we have

$$1 - \left( \frac{\text{Diam}(X) + \rho_i + \rho_j + \sigma\alpha}{\text{Diam}(X) + \sigma\alpha} \right)^\alpha \exp \left( -\frac{\rho_i + \rho_j}{\sigma} \right) \leq \left( \frac{\text{Diam}(X) + \sigma\alpha}{\text{Diam}(X) + \rho_i + \rho_j + \sigma\alpha} \right)^\alpha \exp \left( \frac{\rho_i + \rho_j}{\sigma} \right) - 1,$$

and so Lemma 6 gives the result.

Tighter bounds can be derived if pairwise distances between elements from pairs of clusters are compared directly to the distances between the cluster centers, and for higher dimensional cases the additional tightness can be significant. We prefer to state the result as above due to the sole reliance on the internal cluster radii relative to scaling parameter.

While bounds of the above type are not verifiable for  $L_{\text{norm}}$  due to the fact that it is not diagonally dominant, a similar degree of agreement between the true and approximate eigenvalues has been observed in all cases considered. In this case the  $K \times K$  matrix is given by the normalised Laplacian of the graph of  $P^C(\boldsymbol{\theta})$  with similarities given by  $n_i n_j s(P^C(\boldsymbol{\theta}), i, j)$ . This matrix has the same structure as the original normalised Laplacian, the only difference being the introduction of the factors  $n_i, n_j$ .

Figure 4 shows (a)  $\lambda_2(L(\boldsymbol{\theta}))$  and (b)  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$  plotted against the single projection angle  $\boldsymbol{\theta}$  for the 2 dimensional S1 data set (Fränti and Virtajoki, 2006). The parameter  $\sigma$  was chosen using the same method as for our experiments. A complete linkage clustering was performed for 3000 microclusters (= 60% of total number of data), as well as for 200 microclusters for comparison. The true values of  $\lambda_2(L(\boldsymbol{\theta}))$  and those based on approximations using 3000 microclusters are almost indistinguishable. The approximations based on 200 microclusters also show a good approximation of the optimisation surface, and lie well within the bounds pertaining to the 3000 microcluster case. The same sort of agreement can be seen for  $\lambda_2(L_{\text{norm}}(\boldsymbol{\theta}))$ . Importantly, while the approximations based on 200 microclusters slightly underestimate the true eigenvalues, the location of the local minima, and indeed the shape of the optimisation surface, are very similar to the truth, and so optimising over this approximate surface leads to near optimal projections. We also show the absolute relative error, (c) and (d), as described in Lemma 7. The pessimism of the bound is clearly evident in the bottom left plot where the values of  $\frac{|\lambda_2(L(\boldsymbol{\theta})) - \lambda_2(N(\boldsymbol{\theta}) - B(\boldsymbol{\theta}))|}{\lambda_2(L(\boldsymbol{\theta}))}$  appear very close to zero on the scale of the theoretical bound.

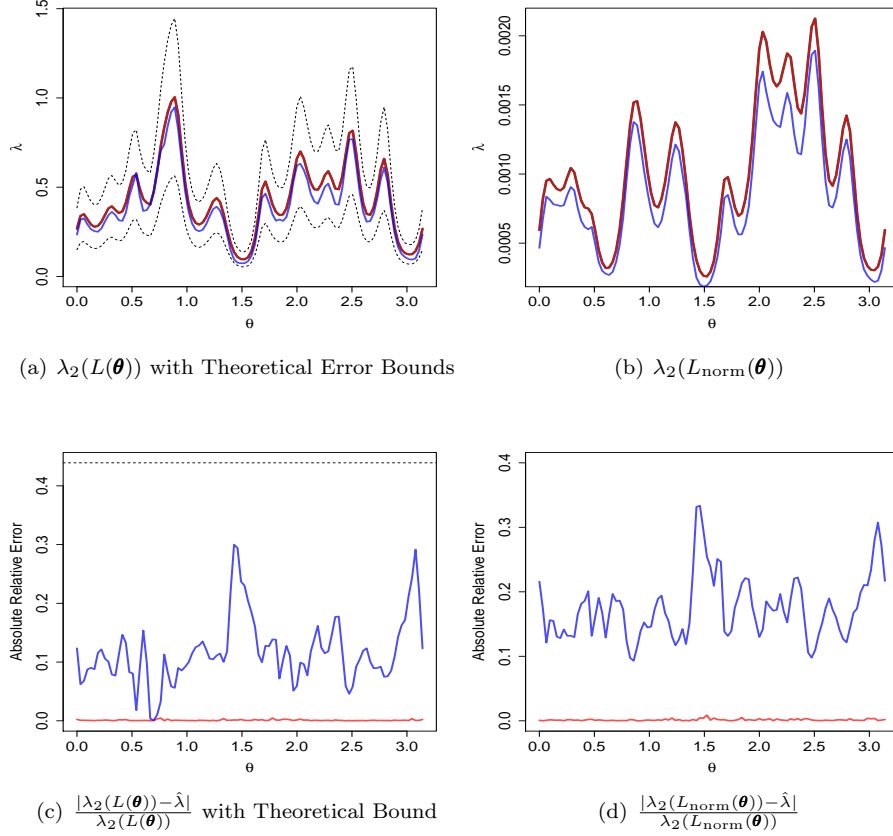
## 7 Experimental Results

In this section we evaluate the proposed method on a large collection of benchmark datasets. We compare our approach with existing dimension reduction methods for clustering, where the final clustering result is determined using spectral clustering. In addition we consider solving our problem iteratively for a shrinking sequence of scaling parameters to find large margin separators, relying on the theoretical results presented in Section 5. We compare these results with the iterative support vector regression approach of Zhang et al. (2009)<sup>3</sup>, a state-of-the-art maximum margin clustering algorithm.

We compare the different methods based on two popular evaluation metrics for clustering, namely purity (Zhao and Karypis, 2004) and  $V$ -measure (Rosenberg and Hirschberg, 2007). Both metrics compare the label assignments made by a clustering algorithm with the true class labels of the data. They take values in  $[0, 1]$  with larger values indicating a better agreement between the two label sets, and hence a superior clustering result. Purity is the weighted average of the largest proportion of

<sup>3</sup>We are grateful to Dr. Kai Zhang for supplying us with code to implement this method.

Figure 4: Approximation Error Plots for S1 data set.



True eigenvalue (—), bounds based on 3000 microclusters (- - -), approximation using 3000 microclusters (—), approximation using 200 microclusters (—)

each cluster which can be represented by a single class.  $V$ -measure is defined as the harmonic mean of measures of completeness and homogeneity. Homogeneity is similar to purity, in that it measures the extent to which each cluster may be represented by a single class, but is given by the weighted average of the entropy of the class distribution within each cluster. Completeness is symmetric to homogeneity, and measures the entropy of the cluster distribution within each class.  $V$ -measure therefore also captures the extent to which classes are split between clusters.

We will use the following notation throughout this section:

- $\text{SCP}^2$  and  $\text{SC}_n\text{P}^2$  refer to the proposed projection pursuits for minimising  $\lambda_2(L(\theta))$  and  $\lambda_2(L_{\text{norm}}(\theta))$  respectively.
- $\text{LMSC}$  refers to the proposed approach of finding large margin separators, based on repeatedly minimising  $\lambda_2(L(\theta))$  for a shrinking sequence of scaling parameters.
- Subscripts “ $o$ ” and “ $c$ ” indicate whether we use an orthogonal projection ( $\omega > 0$ ) or a correlated one ( $\omega < 0$ ), respectively.
- $\text{SC}$  and  $\text{SC}_n$  refer to spectral clustering based on the eigen-decompositions of  $L$  and  $L_{\text{norm}}$  respectively.

- Subscripts “PCA” and “ICA” indicate principal and independent component analysis projections respectively. For example,  $SC_{nPCA}$  refers to spectral clustering using the normalised Laplacian applied to the data projected into a principal component oriented subspace.
- DRSC abbreviates dimensionality reduction for spectral clustering, proposed by Niu et al. (2011). This existing approach applies only to the normalised Laplacian.
- $iSVR_L$  and  $iSVR_G$  denote the iterative support vector regression approach for maximum margin clustering (Zhang et al., 2009), using the linear and Gaussian kernels respectively.

## 7.1 Details of Implementation

To extend our approach to datasets containing multiple clusters, we simply recursively partition (subsets of) the dataset until the desired number of clusters is obtained. We prefer this approach to the alternative of directly seeking a projection which yields a full  $K$  way partition of the dataset, i.e., by minimising the sum of the first  $K$  eigenvalues of the Laplacians, as it is not always clear that all the clusters present in the data can be exposed using a single projection of fixed dimension. At each iteration in this recursive bi-partitioning we split the largest remaining cluster.

The scaling parameter and initialisation are set for each bi-partition, given values determined by the subset of the data being partitioned. For the fixed scaling parameter approaches,  $SCP^2$  and  $SC_nP^2$ , we set  $\sigma = \sqrt{l\lambda_d}N^{-1/5}$ , where  $l$  is the dimension of the projection,  $N$  is the size of the (subset of the) data and  $\lambda_d$  is the largest eigenvalue of the covariance matrix. The value  $\sqrt{\lambda_d}$  captures the scale of the data, while  $\sqrt{l}$  accounts for the fact that distances scale roughly with the square root of the dimension. The denominator term,  $N^{-1/5}$ , is borrowed from kernel density methods and we have found it to work reasonably well for our applications as well. For the large margin approach, LMSC,  $\sigma$  is initialised at  $\sqrt{l\lambda_d}N^{-1/5}$  and decreased by a factor of two with each minimisation of  $\lambda_2(L(\theta))$ , until convergence of the projection matrix. The initialisation of  $V(\theta)$  is via the first  $l$  principal components. For the orthogonal projections we use a two dimensional projection, as this is the lowest dimensional space which can expose non-linear separation between clusters. For the correlated projections we provide a three dimensional initialisation, and it was found that in most cases a high quality univariate projection could be determined from this.

For the LMSC approach, because the values within the Laplacian matrix approach zero, the optimisation becomes less robust, and we found that the correlated approach did not always lead to large margin separation. We believe this is as a result of the term controlling the correlation becoming too dominant relative to the decreasing eigenvalue unless very careful tuning of  $\omega$  is performed. We therefore consider a univariate projection instead of the multivariate correlated approach in this case.

Recall that the parameter  $\beta$  controls the size of the constraint set  $\Delta$ . It is clear that smaller values of  $\beta$  will tend to lead to more balanced partitions, but a precise interpretation of the resulting cluster sizes is unavailable. At best bounds on the cluster sizes can be computed using Chebyshev’s inequality. Rather than relying on these bounds, which may be loose and difficult to interpret in multivariate projections, we recommend applying the proposed method for a range of values of  $\beta$  and selecting the solution corresponding to the largest value of  $\beta$  which induces a specified balance in the partition. We define this balance to be satisfied if the smallest cluster size is at least half the average, i.e.,  $N/2K$ . In this way the effect of the constraint is limited while still producing the desired result. We initialise with a large value of  $\beta$  and decrease by 0.5 until the balance is met. If this balance is not met for  $\beta = 0.5$ , then the corresponding “unbalanced” result is returned anyway.

The parameter  $\delta$  is set to  $\min\{0.01, \sigma^2\}$  and  $\alpha$  to 0.1. We have found these two parameters not to significantly influence the performance of the method. It is important to note, however, that that parameter  $\alpha$  controls the shape of the similarity function, and as a result there is an interplay between this value and the value of  $\sigma$ . For substantially larger values of  $\alpha$  we expect a smaller value of  $\sigma$  to be more appropriate.

For competing approaches based on spectral clustering we do the following. Whenever the number of data exceeds 1000 we use the approximation method of Yan et al. (2009). Following Niu et al.

(2011), we set the reduced dimension to  $K - 1$ , where  $K$  is the number of clusters. We compute clustering results for all values of  $\sigma$  in  $\{0.1, 0.2, 0.5, 1, 2, 5, 10, 20, 50, 100, 200\}$  as well as for the local scaling approach of Zelnik-Manor and Perona (2004), and report the highest performance in each case. For DRSC we also considered the parameter setting used for our method, and to implement the local scalings of Zelnik-Manor and Perona (2004) these were recomputed with each iteration in the corresponding projected subspace. We also provided DRSC with a warm start via PCA as this improved performance over a random initialisation, and offers a more fair comparison. Because of this extensive search over scaling parameters we expect competing methods to achieve very high performance whenever the corresponding dimension reduction is capable of exposing the clusters in the data well.

For the iSVR maximum margin clustering method, we set the balancing parameter equal to 0.3 as suggested by Zhang et al. (2009) when the cluster sizes are not balanced. We argue that the balance of the clusters will not be known in practice, and the unbalanced setting led to superior performance compared with the balanced setting in the examples considered. The iSVR approach also generates only a bi-partition, and to generate multiple clusters we apply the same recursive approach as in our method.

## 7.2 Clustering Results

The following benchmark datasets were used for comparison:

- Optical recognition of handwritten digits (Opt. Digits).<sup>4</sup> 5620  $8 \times 8$  compressed images of handwritten digits in  $\{0, \dots, 9\}$ , resulting in 64 dimensions with 10 classes.
- Pen based recognition of handwritten digits (Pen Digits).<sup>2</sup> 10992 observations, each corresponding to a stylus trajectory ( $x, y$  coordinates) from a handwritten digit in  $\{0, \dots, 9\}$ , i.e., 10 classes. The trajectories are sampled at 8 time points, resulting in 16 dimensions.
- Satellite.<sup>2</sup> 6435 multispectral values from  $3 \times 3$  pixel squares from satellite images, which results in 36 dimensions. There are 6 classes corresponding to different land types.
- Breast cancer Wisconsin (Br. Cancer).<sup>2</sup> 699 observations with 9 attributes relating to tumour masses. There are 2 classes corresponding to benign and malignant masses.
- Congressional votes (Voters).<sup>2</sup> 435 sets of 16 binary decisions on US congressional ballots. The 2 classes correspond to political party membership.
- Dermatology.<sup>2</sup> 366 observations corresponding to dermatology patients, each containing 34 dimensions derived from clinical and histopathological features. There are 6 classes corresponding to different skin diseases.
- Yeast cell cycle analysis (Yeast).<sup>5</sup> 698 yeast genes across 72 conditions (dimensions). There are 5 classes corresponding to different genotypes.
- Synthetic control chart (Chart).<sup>2</sup> 600 simulated time series of length 60 displaying one of 6 fundamental characteristics, leading to 6 classes.
- Multiple feature digits (M.F. Digits).<sup>2</sup> 2000 handwritten digits in  $\{0, \dots, 9\}$  taken from Dutch utility maps. Following Niu et al. (2011) we use only the 216 profile correlation features.
- Statlog image segmentation (Image Seg.).<sup>2</sup> 2310 observations containing 19 features derived from  $3 \times 3$  pixel squares from 7 outdoor images. Each image constitutes a class.

<sup>4</sup><https://archive.ics.uci.edu/ml/datasets.html>

<sup>5</sup><http://genome-www.stanford.edu/cellcycle/>



Table 1: Purity results for spectral clustering using the standard Laplacian,  $L$ . Average performance from 30 runs on each dataset, with standard deviation as subscript. The highest average performance in each case is highlighted in bold.

	$SCP_o^2$	$SCP_c^2$	$SC_{PCA}$	$SC_{ICA}$	$SC$
Opt. Digits	0.81 <sub>0.05</sub>	<b>0.83</b> <sub>0.06</sub>	0.33 <sub>0.05</sub>	0.20 <sub>0.02</sub>	0.11 <sub>0.00</sub>
Pen Digits	<b>0.78</b> <sub>0.01</sub>	0.76 <sub>0.01</sub>	0.64 <sub>0.02</sub>	0.53 <sub>0.03</sub>	0.62 <sub>0.03</sub>
Satellite	<b>0.75</b> <sub>0.00</sub>	0.73 <sub>0.03</sub>	0.63 <sub>0.01</sub>	0.72 <sub>0.01</sub>	0.62 <sub>0.02</sub>
Br. Cancer	<b>0.97</b> <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	0.96 <sub>0.00</sub>
Voters	0.84 <sub>0.00</sub>	0.83 <sub>0.00</sub>	<b>0.86</b> <sub>0.00</sub>	<b>0.86</b> <sub>0.00</sub>	0.78 <sub>0.00</sub>
Dermatology	<b>0.94</b> <sub>0.00</sub>	0.90 <sub>0.00</sub>	0.91 <sub>0.00</sub>	0.89 <sub>0.00</sub>	0.56 <sub>0.00</sub>
Yeast	<b>0.75</b> <sub>0.00</sub>	0.74 <sub>0.00</sub>	0.67 <sub>0.00</sub>	0.62 <sub>0.00</sub>	0.60 <sub>0.00</sub>
Chart	<b>0.84</b> <sub>0.00</sub>	0.83 <sub>0.00</sub>	0.72 <sub>0.06</sub>	0.65 <sub>0.07</sub>	0.56 <sub>0.05</sub>
M.F. Digits	<b>0.83</b> <sub>0.02</sub>	0.79 <sub>0.02</sub>	0.53 <sub>0.03</sub>	0.34 <sub>0.03</sub>	0.31 <sub>0.04</sub>
Image Seg.	0.62 <sub>0.03</sub>	<b>0.68</b> <sub>0.03</sub>	0.53 <sub>0.02</sub>	0.46 <sub>0.02</sub>	0.56 <sub>0.03</sub>

Before applying the clustering algorithms, data were rescaled so that every feature had unit variance. This is a standard approach to handle situations where different features are captured on different scales and an appropriate rescaling is not obviously apparent from the context. For consistency we used this same preprocessing approach for all datasets.

### 7.2.1 Spectral Clustering Using the Standard Laplacian

Tables 1 and 2 report the purity and  $V$ -measure respectively for the proposed method and spectral clustering using the standard Laplacian applied to the original data, as well as their projection into PCA and ICA oriented subspaces. The tables report the average and standard deviation (as subscript) from 30 repetitions. The highest average performance for each dataset is highlighted in bold. Both the orthogonal and correlated projection approaches achieve substantially higher performance than other methods in the majority of cases. There are few cases where they are not competitive with the best performing of the competing approaches, while there are multiple examples where the proposed methods strongly outperform all others. The two versions of the proposed method are closely comparable with one another on average, with the correlated approach offering a slightly better worst case comparison. This, however, does not appear highly significant beyond sampling variation both within each dataset and with respect to the collection of datasets used for comparison. What is evident is that the added flexibility offered by multivariate projections does not result in a substantial improvement over univariate projections, which induce linear cluster boundaries.

### 7.2.2 Spectral Clustering Using the Normalised Laplacian

Tables 3 and 4 report the purity and  $V$ -measure respectively for the proposed approach based on minimising  $\lambda_2(L_{\text{norm}}(\theta))$ , the dimensionality reduction for spectral clustering algorithm (Niu et al., 2011) and spectral clustering based on the normalised Laplacian applied to the original data and their PCA and ICA projections. Again the tables show the average and standard deviation from 30 runs of each method, with the highest average performance on each dataset highlighted in bold.

The proposed approach using both correlated and orthogonal projections is again competitive with all other methods in almost all cases considered. In addition both versions of the proposed approach substantially outperform the other methods in multiple examples. Unlike in the case of the standard Laplacian, here there is evidence that the orthogonal projection achieves better clustering results in general, outperforming the correlated approach in the majority of examples.

Table 2:  $V$ -measure results for spectral clustering using the standard Laplacian,  $L$ . Average performance from 30 runs on each dataset, with standard deviation as subscript. The highest average performance in each case is highlighted in bold.

	$SCP_o^2$	$SCP_c^2$	$SC_{PCA}$	$SC_{ICA}$	$SC$
Opt. Digits	0.77 <sub>0.03</sub>	<b>0.78</b> <sub>0.04</sub>	0.40 <sub>0.05</sub>	0.21 <sub>0.03</sub>	0.01 <sub>0.00</sub>
Pen Digits	<b>0.75</b> <sub>0.01</sub>	0.74 <sub>0.02</sub>	0.66 <sub>0.01</sub>	0.54 <sub>0.03</sub>	0.64 <sub>0.02</sub>
Satellite	<b>0.60</b> <sub>0.00</sub>	0.59 <sub>0.03</sub>	0.50 <sub>0.00</sub>	<b>0.60</b> <sub>0.01</sub>	0.50 <sub>0.01</sub>
Br. Cancer	0.79 <sub>0.00</sub>	<b>0.80</b> <sub>0.00</sub>	0.78 <sub>0.00</sub>	0.78 <sub>0.00</sub>	0.77 <sub>0.00</sub>
Voters	<b>0.42</b> <sub>0.00</sub>	0.38 <sub>0.00</sub>	0.41 <sub>0.00</sub>	0.41 <sub>0.00</sub>	0.30 <sub>0.00</sub>
Dermatology	<b>0.89</b> <sub>0.00</sub>	0.83 <sub>0.00</sub>	0.86 <sub>0.00</sub>	0.82 <sub>0.00</sub>	0.58 <sub>0.00</sub>
Yeast	<b>0.54</b> <sub>0.00</sub>	<b>0.54</b> <sub>0.00</sub>	0.51 <sub>0.00</sub>	0.40 <sub>0.00</sub>	0.41 <sub>0.00</sub>
Chart	0.77 <sub>0.00</sub>	0.77 <sub>0.00</sub>	<b>0.81</b> <sub>0.02</sub>	0.73 <sub>0.04</sub>	0.70 <sub>0.02</sub>
M.F. Digits	<b>0.76</b> <sub>0.01</sub>	0.73 <sub>0.02</sub>	0.56 <sub>0.02</sub>	0.36 <sub>0.05</sub>	0.38 <sub>0.05</sub>
Image Seg.	0.62 <sub>0.01</sub>	<b>0.65</b> <sub>0.02</sub>	0.53 <sub>0.01</sub>	0.46 <sub>0.01</sub>	0.58 <sub>0.02</sub>

Table 3: Purity results for spectral clustering using the normalised Laplacian,  $L_{\text{norm}}$ . Average performance from 30 runs on each dataset, with standard deviation as subscript. The highest average performance in each case is highlighted in bold.

	$SC_n P_o^2$	$SC_n P_c^2$	DRSC	$SC_{nPCA}$	$SC_{nICA}$	$SC_n$
Opt. Digits	<b>0.81</b> <sub>0.05</sub>	0.74 <sub>0.06</sub>	0.80 <sub>0.03</sub>	0.66 <sub>0.03</sub>	0.64 <sub>0.01</sub>	0.66 <sub>0.02</sub>
Pen Digits	<b>0.78</b> <sub>0.00</sub>	0.74 <sub>0.02</sub>	0.69 <sub>0.04</sub>	0.76 <sub>0.03</sub>	0.74 <sub>0.01</sub>	0.77 <sub>0.04</sub>
Satellite	0.75 <sub>0.00</sub>	0.74 <sub>0.03</sub>	0.73 <sub>0.01</sub>	<b>0.76</b> <sub>0.01</sub>	0.73 <sub>0.02</sub>	0.74 <sub>0.01</sub>
Br. Cancer	<b>0.97</b> <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	0.96 <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>
Voters	0.85 <sub>0.00</sub>	0.84 <sub>0.00</sub>	<b>0.86</b> <sub>0.00</sub>	<b>0.86</b> <sub>0.00</sub>	<b>0.86</b> <sub>0.00</sub>	0.85 <sub>0.00</sub>
Dermatology	0.86 <sub>0.00</sub>	0.91 <sub>0.00</sub>	0.87 <sub>0.00</sub>	0.92 <sub>0.02</sub>	0.91 <sub>0.00</sub>	<b>0.95</b> <sub>0.00</sub>
Yeast	<b>0.76</b> <sub>0.00</sub>	0.70 <sub>0.00</sub>	0.62 <sub>0.00</sub>	0.71 <sub>0.00</sub>	0.69 <sub>0.01</sub>	0.60 <sub>0.00</sub>
Chart	<b>0.87</b> <sub>0.00</sub>	0.85 <sub>0.00</sub>	0.75 <sub>0.00</sub>	0.71 <sub>0.06</sub>	0.80 <sub>0.07</sub>	0.75 <sub>0.00</sub>
M.F. Digits	<b>0.84</b> <sub>0.01</sub>	0.79 <sub>0.01</sub>	0.77 <sub>0.03</sub>	0.77 <sub>0.03</sub>	0.79 <sub>0.01</sub>	0.77 <sub>0.02</sub>
Image Seg.	0.66 <sub>0.01</sub>	<b>0.70</b> <sub>0.01</sub>	0.65 <sub>0.04</sub>	0.61 <sub>0.03</sub>	0.55 <sub>0.02</sub>	0.62 <sub>0.02</sub>

Table 4:  $V$ -measure results for spectral clustering using the normalised Laplacian,  $L_{\text{norm}}$ . Average performance from 30 runs on each dataset, with standard deviation as subscript. The highest average performance in each case is highlighted in bold.

	$SC_n P_o^2$	$SC_n P_c^2$	DRSC	$SC_{nPCA}$	$SC_{nICA}$	$SC_n$
Opt. Digits	<b>0.77</b> <sub>0.03</sub>	0.72 <sub>0.04</sub>	0.75 <sub>0.02</sub>	0.62 <sub>0.02</sub>	0.60 <sub>0.01</sub>	0.62 <sub>0.01</sub>
Pen Digits	<b>0.76</b> <sub>0.00</sub>	0.74 <sub>0.01</sub>	0.66 <sub>0.02</sub>	0.72 <sub>0.01</sub>	0.68 <sub>0.01</sub>	0.73 <sub>0.02</sub>
Satellite	0.61 <sub>0.01</sub>	0.60 <sub>0.03</sub>	0.57 <sub>0.01</sub>	<b>0.62</b> <sub>0.01</sub>	0.60 <sub>0.05</sub>	0.60 <sub>0.01</sub>
Br. Cancer	0.79 <sub>0.00</sub>	<b>0.81</b> <sub>0.00</sub>	0.76 <sub>0.00</sub>	<b>0.81</b> <sub>0.00</sub>	<b>0.81</b> <sub>0.00</sub>	0.79 <sub>0.00</sub>
Voters	<b>0.42</b> <sub>0.00</sub>	0.41 <sub>0.00</sub>	<b>0.42</b> <sub>0.00</sub>	0.41 <sub>0.00</sub>	0.41 <sub>0.00</sub>	<b>0.42</b> <sub>0.00</sub>
Dermatology	0.85 <sub>0.00</sub>	0.86 <sub>0.00</sub>	0.80 <sub>0.00</sub>	0.86 <sub>0.02</sub>	0.83 <sub>0.00</sub>	<b>0.91</b> <sub>0.00</sub>
Yeast	<b>0.57</b> <sub>0.00</sub>	0.45 <sub>0.00</sub>	0.41 <sub>0.00</sub>	0.54 <sub>0.00</sub>	0.53 <sub>0.00</sub>	0.45 <sub>0.00</sub>
Chart	<b>0.81</b> <sub>0.00</sub>	0.80 <sub>0.00</sub>	0.75 <sub>0.00</sub>	<b>0.81</b> <sub>0.02</sub>	0.80 <sub>0.05</sub>	0.73 <sub>0.00</sub>
M.F. Digits	<b>0.76</b> <sub>0.01</sub>	0.74 <sub>0.02</sub>	0.69 <sub>0.02</sub>	0.71 <sub>0.02</sub>	0.75 <sub>0.01</sub>	0.70 <sub>0.02</sub>
Image Seg.	0.65 <sub>0.01</sub>	<b>0.68</b> <sub>0.01</sub>	0.62 <sub>0.04</sub>	0.60 <sub>0.02</sub>	0.46 <sub>0.02</sub>	0.60 <sub>0.02</sub>

Table 5: Purity results for large margin clustering. Average performance from 30 runs on each dataset, with standard deviation as subscript. The highest average performance in each case is highlighted in bold.

	LMSC <sub>o</sub>	LMSC	iSVR <sub>L</sub>	iSVR <sub>G</sub>
Opt. Digits	0.74 <sub>0.05</sub>	0.69 <sub>0.05</sub>	<b>0.76</b> <sub>0.00</sub>	0.61 <sub>0.01</sub>
Pen Digits	<b>0.80</b> <sub>0.02</sub>	0.71 <sub>0.04</sub>	0.78 <sub>0.00</sub>	0.78 <sub>0.00</sub>
Satellite	<b>0.75</b> <sub>0.01</sub>	0.72 <sub>0.03</sub>	0.68 <sub>0.00</sub>	0.68 <sub>0.00</sub>
Br. Cancer	0.96 <sub>0.00</sub>	<b>0.97</b> <sub>0.00</sub>	0.90 <sub>0.00</sub>	0.95 <sub>0.00</sub>
Voters	<b>0.84</b> <sub>0.00</sub>	<b>0.84</b> <sub>0.00</sub>	0.81 <sub>0.00</sub>	<b>0.84</b> <sub>0.00</sub>
Dermatology	<b>0.86</b> <sub>0.00</sub>	<b>0.86</b> <sub>0.00</sub>	0.80 <sub>0.00</sub>	0.80 <sub>0.00</sub>
Yeast	<b>0.75</b> <sub>0.00</sub>	<b>0.75</b> <sub>0.00</sub>	<b>0.75</b> <sub>0.00</sub>	0.70 <sub>0.00</sub>
Chart	<b>0.89</b> <sub>0.00</sub>	0.83 <sub>0.00</sub>	0.72 <sub>0.00</sub>	0.72 <sub>0.00</sub>
M.F. Digits	<b>0.82</b> <sub>0.05</sub>	0.74 <sub>0.04</sub>	0.67 <sub>0.00</sub>	0.60 <sub>0.01</sub>
Image Seg.	0.60 <sub>0.04</sub>	0.65 <sub>0.03</sub>	0.64 <sub>0.00</sub>	<b>0.66</b> <sub>0.01</sub>

### 7.2.3 Large Margin Clustering

It is important to note that the method described in Section 6 does not provide a close approximation as  $\sigma \rightarrow 0^+$ . For the datasets containing more than 1000 data we use the microcluster approach for all values of  $\sigma$  and therefore only guarantee a large separation between the microclusters. It is arguable that this is a preferable objective as the maximum margin is not robust in the presence of noise, and it is not clear that it converges in the general setting (Ben-David et al., 2009). Microclusters have the potential to absorb some of the noise in the data, and in the event that they are of roughly equal density, maximising the margin over the microcluster centers has a similar effect to that of minimising the empirical density in a neighbourhood of the corresponding hyperplane separator. This is reminiscent of the soft-margin approach which does enjoy strong convergence properties (Ben-David et al., 2009). In addition, since the optimisation is reinitialised for each value of  $\sigma$ , we are able to recompute the microclusters by performing the coarse clustering on the projected data with each iteration. This tends to lead to the margins in the microclusters being more closely related to the margins in the full dataset along the optimal projections.

Tables 5 and 6 report the average and standard deviation of the proposed LMSC as well as the iterative support vector regression approach (Zhang et al., 2009) using both a linear and Gaussian kernel for comparison. Both versions of LMSC, using an orthogonal two dimensional and a one dimensional projection, outperform both versions of the iterative support vector regression in the majority of cases, with substantially higher performance in multiple examples. There is strong evidence that the two dimensional LMSC<sub>o</sub> obtains better quality clustering results than the one dimensional alternative, showing substantially higher performance in the vast majority of cases considered.

## 7.3 Summarising Clustering Performance

Thus far we have compared different approaches for standard and normalised spectral clustering and for large margin clustering separately. These separate comparisons are important to understand the benefits of the proposed methods, however when considering the clustering problem abstractly it is necessary to compare all methods jointly. It is already clear that no method is uniformly superior to all others, since even within the separate comparisons no method outperformed the rest in every example. We find it important to reiterate the fact that for competing methods based on spectral clustering an extensive search over scaling parameters was performed and the best performance reported, whereas for our method only a simple data driven heuristic was used in every example. Such a search is not possible in practice since the true labels will not be known, and hence the results reported for these methods likely overestimate their true expected performance in practice. What was evident, however,

Table 6:  $V$ -measure results for large margin clustering. Average performance from 30 runs on each dataset, with standard deviation as subscript. The highest average performance in each case is highlighted in bold.

	LMSC <sub>o</sub>	LMSC	iSVR <sub>L</sub>	iSVR <sub>G</sub>
Opt. Digits	0.70 <sub>0.03</sub>	0.62 <sub>0.04</sub>	<b>0.72</b> <sub>0.00</sub>	0.57 <sub>0.00</sub>
Pen Digits	<b>0.76</b> <sub>0.02</sub>	0.66 <sub>0.03</sub>	0.72 <sub>0.01</sub>	<b>0.73</b> <sub>0.00</sub>
Satellite	<b>0.61</b> <sub>0.01</sub>	0.57 <sub>0.03</sub>	0.55 <sub>0.00</sub>	0.55 <sub>0.00</sub>
Br. Cancer	0.76 <sub>0.00</sub>	<b>0.78</b> <sub>0.00</sub>	0.55 <sub>0.00</sub>	0.72 <sub>0.00</sub>
Voters	<b>0.43</b> <sub>0.00</sub>	0.38 <sub>0.00</sub>	0.34 <sub>0.00</sub>	0.42 <sub>0.00</sub>
Dermatology	<b>0.86</b> <sub>0.00</sub>	0.85 <sub>0.00</sub>	0.77 <sub>0.00</sub>	0.74 <sub>0.01</sub>
Yeast	0.56 <sub>0.00</sub>	<b>0.58</b> <sub>0.00</sub>	0.55 <sub>0.01</sub>	0.53 <sub>0.00</sub>
Chart	<b>0.85</b> <sub>0.00</sub>	0.78 <sub>0.00</sub>	0.66 <sub>0.00</sub>	0.72 <sub>0.00</sub>
M.F. Digits	<b>0.75</b> <sub>0.03</sub>	0.69 <sub>0.03</sub>	0.60 <sub>0.00</sub>	0.62 <sub>0.01</sub>
Image Seg.	0.60 <sub>0.03</sub>	<b>0.63</b> <sub>0.03</sub>	<b>0.63</b> <sub>0.00</sub>	0.60 <sub>0.01</sub>

is that the local scaling approach of Zelnik-Manor and Perona (2004) is very effective and yielded the highest performance in roughly half the cases considered.

It is clearly apparent from the performance of the various methods that the clustering problem differs vastly in difficulty across the different datasets considered. To combine the results from the different datasets we standardise them as follows. For each dataset  $D$  we compute for each method the relative deviation from the average performance of all methods when applied to  $D$ . That is, for each method  $M_i$  we compute the relative purity,

$$\text{Rel.Purity}(M_i, D) = \frac{\text{Purity}(M_i, D) - \frac{1}{\#\text{Methods}} \sum_{j=1}^{\#\text{Methods}} \text{Purity}(M_j, D)}{\frac{1}{\#\text{Methods}} \sum_{j=1}^{\#\text{Methods}} \text{Purity}(M_j, D)}, \quad (31)$$


and similarly for  $V$ -measure. We can then compare the distributions of the relative performance measures from all datasets and for all methods. It is clear from Table 1 that the competing methods SC, SC<sub>PCA</sub> and SC<sub>ICA</sub> are not competitive with other methods in general, due to their vastly inferior performance on multiple datasets. Moreover, their performance is sufficiently low to obscure the comparisons between others. These three methods are therefore omitted from this comparison. Figures 5 and 6 show boxplots of the relative performance measures for all other methods. The additional red dots indicate the mean relative performance measures for each method, and methods are ordered in decreasing order of their means. In the case of purity, all of the proposed methods outperform every method used for comparison, and except for the univariate large margin method, LMSC, the difference between the proposed methods and the methods used for comparison is substantial. In the case of  $V$ -measure, the same is true except that in this case LMSC is outperformed on average by spectral clustering using the normalised Laplacian applied to the PCA projected data. Notice that the most relevant comparison for LMSC is with iSVR<sub>L</sub> because of their similar objectives. In terms of both purity and  $V$ -measure, LMSC significantly outperformed the existing large margin clustering method.

Among the methods used for comparison, it is evident that spectral clustering is capable of outperforming existing large margin clustering methods, provided an appropriate scaling parameter can be determined. Of those spectral clustering variants, PCA projections showed the best overall performance. While the DRSC method (Niu et al., 2011) in some cases showed a substantial improvement over the simpler dimension reduction of PCA, it did not yield consistently higher performance on the datasets considered.

Overall it is apparent that the proposed approach for projection pursuit based on spectral connectivity is highly competitive with existing dimension reduction methods. Moreover, a simple data driven heuristic allowed us to select the important scaling parameter automatically without tuning it for each dataset, as is recommended for the DRSC method (Niu et al., 2011). Among the variants of

the proposed approaches, it is evident that while the flexibility of the multivariate projections offered higher performance on average than the corresponding univariate projections, it is only in the case of the large margin separation methods that this improvement is significant beyond the variation from the collection of datasets used for comparison.

Figure 5: Box plots of relative purity with additional red dots to indicate means. Methods are ordered with decreasing mean value.




## 7.4 Sensitivity Analysis

In this subsection we investigate the sensitivity of the proposed approach to the setting of the important scaling parameter,  $\sigma$ . In addition we consider the effect on performance of the number of microclusters used in approximating the optimisation surface. For the former we consider the breast cancer, voters, dermatology, yeast and chart datasets as these exhibited very low variability in performance and offer more interpretable comparisons. Figures 7 and 8 show plots of the purity and  $V$ -measure values for  $\sigma$  taking values in  $\{0.1\sigma_0, 0.2\sigma_0, 0.5\sigma_0, \sigma_0, 2\sigma_0, 5\sigma_0, 10\sigma_0\}$ , where  $\sigma_0 = \sqrt{l\lambda_d N^{-1/5}}$  is the value used in the experiments above. There is some variability in the performance for different values, with no clear pattern indicating that a higher or lower value than the one used is better in general. Importantly there are very few occurrences of substantially poorer performance than that obtained with our simply chosen heuristic, and also it is clear that in the majority of cases performance could be improved from what is reported above if an appropriate tuning of  $\sigma$  is possible.

To investigate the effect of microclusters on clustering accuracy we simulated datasets from Gaussian mixtures containing 5 components (clusters) in 50 dimensions. This allows us to generate datasets of any desired size. For these experiments 30 sets of parameters for the Gaussian mixtures were generated randomly. In the first case a single dataset of size 1000 was simulated from each set of parameters, and clustering solutions obtained for a number of microclusters,  $K$ , ranging from 100 to 1000, the final value therefore applying no approximation. Figure 9 shows the median and interquartile range of both performance measures for 10 values of  $K$ . It is evident that aside from  $K = 100$ , performance is similar for all other values, and so using a small value, say  $K = 200$ , should be sufficient to obtain a good approximation of the underlying optimisation surface.

Figure 6: Box plots of relative  $V$ -measure with additional red dots to indicate means. Methods are ordered with decreasing mean value.



In the second, we fix the number of microclusters,  $K = 200$ , and for each set of parameters simulate datasets with between 1000 and 10 000 observations. In the most extreme case, therefore, the number of microclusters is only 2% of the total number of data. Figure 10 shows the corresponding performance plots, again containing the medians and interquartile ranges. Even for datasets of size 10 000, the coarse approximation of the dataset through 200 microclusters is sufficient to obtain a high quality projection using the proposed approach.


## 8 Conclusions

We proposed a projection pursuit method for finding the optimal subspace in which to perform a binary partition of unlabelled data. The proposed method optimises the separability of the projected data, as measured by spectral graph theory, by minimising the second smallest eigenvalue of the graph Laplacians. The Lipschitz continuity and differentiability properties of this projection index with respect to the projection matrix were established, which enabled us to apply a generalised gradient descent method to find locally optimal solutions. Compared with existing dimension reduction for spectral clustering, we derive expressions for the gradient of the overall objective and so find solutions within a single generalised gradient descent scheme, with guaranteed convergence to a local optimum. Our experiments suggest that the proposed method substantially outperforms spectral clustering applied to the original data as well as existing dimensionality reduction methods for spectral clustering.


A connection to maximal margin hyperplanes was established, showing that in the univariate case, as the scaling parameter of the similarity function is reduced towards zero, the binary partition of the projected data maximises the linear separability between the two clusters. Implementing our method for a shrinking sequence of scaling parameters thus allows us to find large margin separators practically. We found that this approach outperforms state of the art methods for maximum margin clustering on a large collection of datasets.

The computational cost of the proposed projection pursuit method per iteration is  $\mathcal{O}(dN^2)$ , where  $N$  is the number of observations, and  $d$  is the dimensionality, which can become prohibitive for

Figure 7: Sensitivity analysis for varying  $\sigma$ . Standard Laplacian. The  $x$ -axis contains the multiplication factor applied to the default scaling parameter used in the experiments.




(a)  $SCP_o^2$




(b)  $SCP_c^2$

Br. Cancer (+), Voters ( $\triangle$ ), Dermatology ( $\times$ ), Yeast ( $\circ$ ), Chart ( $\diamond$ )

Figure 8: Sensitivity analysis for varying  $\sigma$ . Normalised Laplacian. The  $x$ -axis contains the multiplication factor applied to the default scaling parameter used in the experiments.



(a)  $SC_n P_o^2$



(b)  $SC_n P_c^2$

Br. Cancer (+), Voters ( $\triangle$ ), Dermatology ( $\times$ ), Yeast ( $\circ$ ), Chart ( $\diamond$ )



Figure 9: Sensitivity analysis for varying number of microclusters,  $K$ . Plots show median and interquartile ranges of performance measures from 30 datasets simulated from 50 dimensional Gaussian mixtures with 5 clusters and 1000 observations.

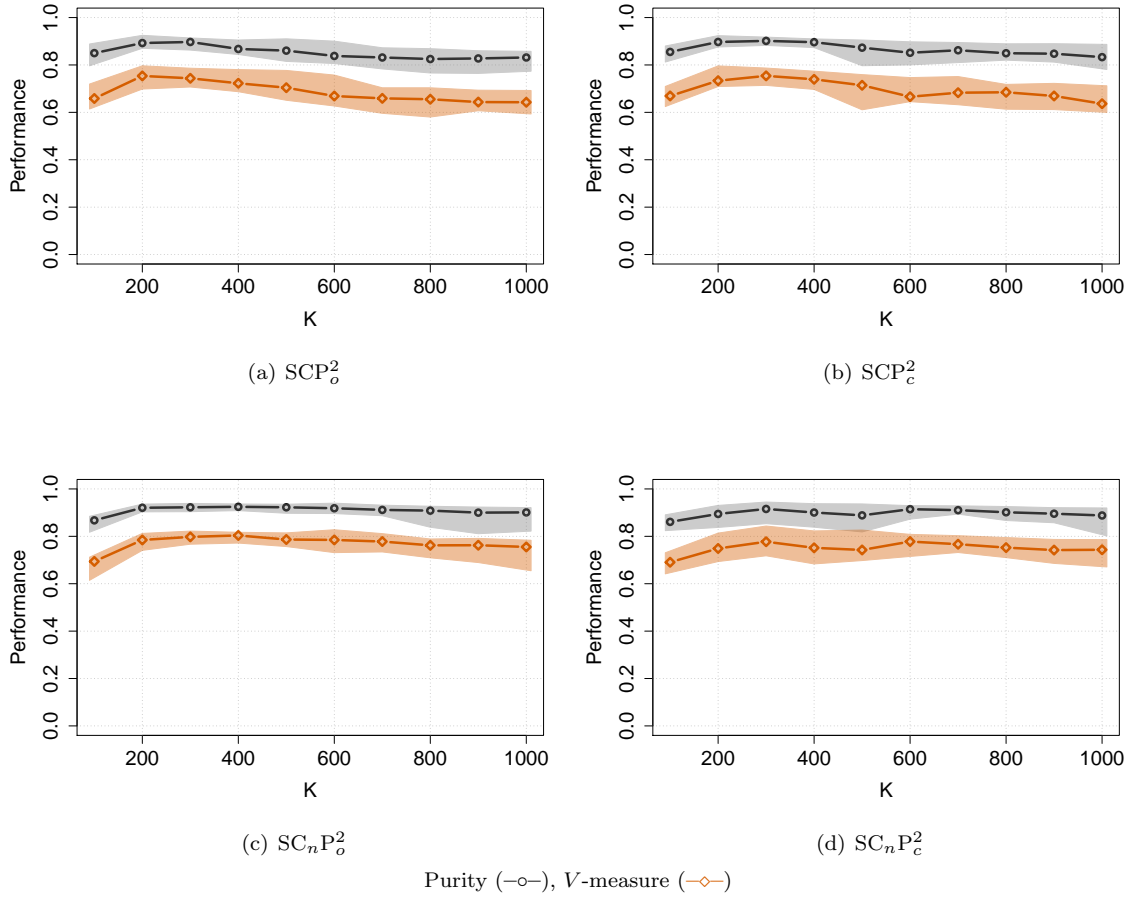
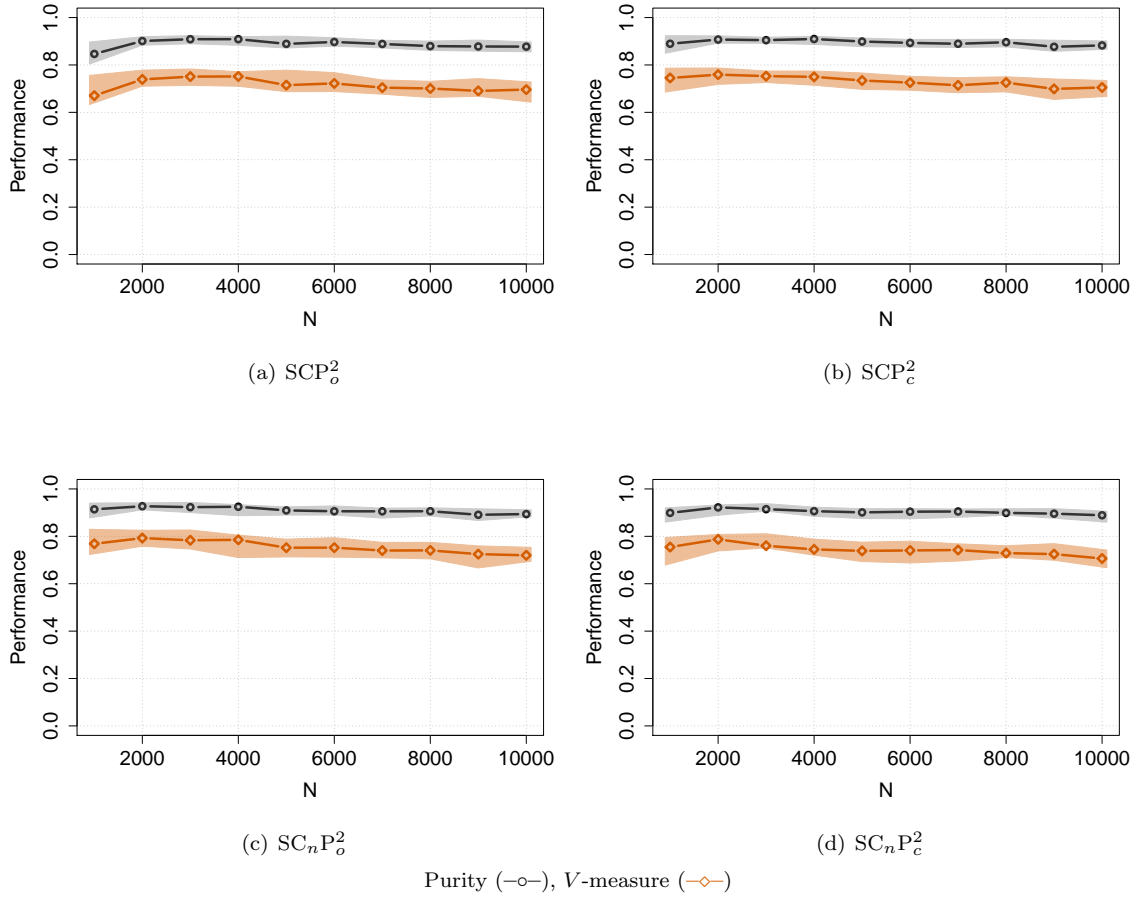


Figure 10: Sensitivity analysis for fixed number of microclusters,  $K = 200$ , and varying number of data. Plots show median and interquartile ranges of performance measures from datasets simulated from 50 dimensional Gaussian mixtures with 5 clusters and between 1000 and 10 000 observations.



large datasets. To ameliorate this an approximation method using microclusters, with provable error bounds is proposed. Our sensitivity analysis, based on clustering performance, indicates that even for relatively few microclusters, the approximation of the optimisation surface is adequate for finding high quality subspaces for clustering.

## A. Derivatives

In the general case we may consider a set of  $K$  microclusters with centers  $c_1, \dots, c_K$  and counts  $n_1, \dots, n_K$ . The derivations we provide in this appendix are valid for  $n_i = 1 \forall i \in \{1, \dots, K\}$ , and so apply to the exact formulation of the problem as well. Let  $\theta \in \Theta$  and let  $P$  be the repeated projected cluster centers,  $P = \{p_1, \dots, p_K, p_K\} = \{V(\theta)^\top c_1, V(\theta)^\top c_1, \dots, V(\theta)^\top c_K\}$ , where each  $V(\theta)^\top c_i$  is repeated  $n_i$  times. In Section 4 we expressed  $D_\theta \lambda$  via the chain rule decomposition  $D_P \lambda D_v P D_\theta v$ . The compression of  $P$  to the size  $K$  non-repeated projected set,  $P^C = \{p_1, \dots, p_K\}$ , requires a slight restructuring, as described in Section 6.

We begin with the standard Laplacian, and define  $N(\theta)$  and  $B(\theta)$  as in Lemma 6. That is,  $N(\theta)$  is the diagonal matrix with  $i$ -th diagonal element equal to  $\sum_{j=1}^K n_j s(P^C, i, j)$  and  $B(\theta)_{i,j} = \sqrt{n_i n_j} s(P^C, i, j)$ . The derivative of the second eigenvalue of the Laplacian of  $P$  relies on the corresponding eigenvector,  $u$ . However, this vector is not explicitly available as we only solve the  $K \times K$  eigen-problem of  $N(\theta) - B(\theta)$ . Let  $u^C$  be the second eigenvector of  $N(\theta) - B(\theta)$ . As in the proof of Lemma 6 if  $i, j$  are such that the  $i$ -th element of  $P$  corresponds to the  $j$ -th microcluster, then  $u_j^C = \sqrt{n_j} u_i$ . The derivative of  $\lambda_2(N(\theta) - B(\theta))$  with respect to the  $i$ -th column of  $\theta$ , and thus equivalently of the second eigenvalue of the Laplacian of  $P$ , is therefore given by

$$\frac{1}{2} \left( \sum_{j,k} \left( \frac{u_j^C}{\sqrt{n_j}} - \frac{u_k^C}{\sqrt{n_k}} \right)^2 n_j n_k \frac{\partial s(P^C, j, k)}{\partial P_{i1}} \dots \sum_{j,k} \left( \frac{u_j^C}{\sqrt{n_j}} - \frac{u_k^C}{\sqrt{n_k}} \right)^2 n_j n_k \frac{\partial s(P^C, j, k)}{\partial P_{iK}} \right) \begin{pmatrix} c_1 & \dots & c_K \end{pmatrix}^\top D_{\theta_i} V_i, \quad (32)$$

where  $(c_1 \dots c_K)$  is the matrix with  $i$ -th column  $c_i$ ,  $P$  is treated as a  $l \times N$  matrix with  $i$ -th column  $p_i$ , and  $D_{\theta_i} V_i$  is given in Eq. (17). Now, the use of the constraint set  $\Delta_\theta$  and the associated transformation makes a further decomposition convenient. Let  $T = \{t_1, \dots, t_K\} = \{T_{\Delta_\theta}(p_1), \dots, T_{\Delta_\theta}(p_K)\}$ . We provide expressions for the specific constraint sets used, i.e.,  $\Delta_\theta = \prod_{i=1}^l [\mu_{\theta_i} - \beta \sigma_{\theta_i}, \mu_{\theta_i} + \beta \sigma_{\theta_i}]$ , where  $\mu_{\theta_i} = \frac{1}{N} \sum_{j=1}^K n_j P_{ij}$  and  $\sigma_{\theta_i}$  is approximated by  $\sqrt{\frac{1}{N} \sum_{j=1}^K n_j (P_{ij} - \mu_{\theta_i})^2}$ . For ease of exposition we assume that each  $\mu_{\theta_i}$  is equal to zero, noting that no generality is lost through this simplification since the value of the eigenvalue of the Laplacian is location independent. The data can therefore be centered prior to projection pursuit and the following formulation employed. We can then express the first component of (32) as  $D_{T_i} \lambda D_{P_i^C} T_i$ , where

$$D_{T_i} \lambda = \frac{1}{2} \left( \sum_{j,k} \left( \frac{u_j^C}{\sqrt{n_j}} - \frac{u_k^C}{\sqrt{n_k}} \right)^2 n_j n_k \frac{\partial k\left(\frac{\|t_j - t_k\|}{\sigma}\right)}{\partial T_{i1}} \dots \sum_{j,k} \left( \frac{u_j^C}{\sqrt{n_j}} - \frac{u_k^C}{\sqrt{n_k}} \right)^2 n_j n_k \frac{\partial k\left(\frac{\|t_j - t_k\|}{\sigma}\right)}{\partial T_{iK}} \right) \quad (33)$$

and  $D_{P_i^C} T_i$  is the  $K \times K$  matrix with

$$(D_{P_i^C} T_i)_{j \neq k} = \begin{cases} \frac{\delta(1-\delta)\beta n_k P_{ik}/N\sigma_{\theta_i}}{(-\beta\sigma_{\theta_i} - P_{ij} + (\delta(1-\delta))^{1/\delta})^\delta}, & P_{ij} < -\beta\sigma_{\theta_i} \\ \frac{\beta n_k P_{ik}}{N\sigma_{\theta_i}}, & -\beta\sigma_{\theta_i} \leq P_{ij} \leq \beta\sigma_{\theta_i} \\ \frac{2\beta n_k P_{ik}}{N\sigma_{\theta_i}} - \frac{\delta(1-\delta)\beta n_k P_{ik}/N\sigma_{\theta_i}}{(P_{ij} - \beta\sigma_{\theta_i} + (\delta(1-\delta))^{1/\delta})^\delta}, & P_{ij} > \beta\sigma_{\theta_i} \end{cases} \quad (34)$$

$$(D_{P_i^C} T_i)_{jj} = \begin{cases} \frac{\delta(1-\delta)(1+\beta n_j P_{ij}/N\sigma_{\theta_i})}{(-\beta\sigma_{\theta_i} - P_{ij} + (\delta(1-\delta))^{1/\delta})^\delta}, & P_{ij} < -\beta\sigma_{\theta_i} \\ 1 + \frac{\beta n_j P_{ij}}{N\sigma_{\theta_i}}, & -\beta\sigma_{\theta_i} \leq P_{ij} \leq \beta\sigma_{\theta_i} \\ \frac{2\beta n_j P_{ij}}{N\sigma_{\theta_i}} + \frac{\delta(1-\delta)(1-\beta n_j P_{ij}/N\sigma_{\theta_i})}{(P_{ij} - \beta\sigma_{\theta_i} + (\delta(1-\delta))^{1/\delta})^\delta}, & P_{ij} > \beta\sigma_{\theta_i}. \end{cases} \quad (35)$$

In the above we have used the lower case  $t_j$  to denote the  $j$ -th element of the transformed projected dataset, where the upper case  $T_{ij}$  denotes the  $ij$ -th element of the  $l \times N$  matrix with  $j$ -th column equal to  $t_j$ . The benefit of this further decomposition lies in the fact that the majority of terms in the sums in (33) are zero. In fact,

$$\frac{1}{2} \sum_{j,k} \left( \frac{u_j^C}{\sqrt{n_j}} - \frac{u_k^C}{\sqrt{n_k}} \right) n_j n_k \frac{\partial k \left( \frac{\|t_j - t_k\|}{\sigma} \right)}{\partial T_{im}} = \sum_{j \neq m} \left( \frac{u_j^C}{\sqrt{n_j}} - \frac{u_m^C}{\sqrt{n_m}} \right) n_j n_m \frac{\partial k \left( \frac{\|t_j - t_m\|}{\sigma} \right)}{\partial T_{im}}, \quad (36)$$

where for the function given in Eq. (30) we have,

$$\frac{\partial k \left( \frac{\|t_j - t_m\|}{\sigma} \right)}{\partial T_{im}} = \frac{T_{ij} - T_{im}}{\sigma^2 \alpha} \left( \frac{\|t_j - t_m\|}{\sigma \alpha} + 1 \right)^{\alpha-1} \exp \left( \frac{\|t_j - t_m\|}{\sigma} \right). \quad (37)$$

For the normalised Laplacian, the reduced  $K \times K$  eigenproblem has precisely the same form as the original  $N \times N$  problem, with the only difference being the introduction of the factors  $n_j n_k$ . In particular, the second eigenvalue of the normalised Laplacian of  $P$  is equal to the second eigenvalue of the Laplacian of the graph of  $P^C$  with similarities given by  $n_j n_k s(P^C, j, k)$ . With the derivation in Section 4 we can see that the corresponding derivative is as for the standard Laplacian, except that the coefficients  $(u_j^C/\sqrt{n_j} - u_k^C/\sqrt{n_k})^2 n_j n_k$  in Eq. (36) are replaced with  $(u_j^C/\sqrt{d_j} - u_k^C/\sqrt{d_k})^2 - \lambda((u_j^C)^2/d_j + (u_k^C)^2/d_k)$ , where  $\lambda$  is the second eigenvalue of the normalised Laplacian of  $P^C$ ,  $u^C$  is the corresponding eigenvector and  $d_j$  is the degree of the  $j$ -th element of  $P^C$ .

\*Bibliography

- S. Ben-David, T. Lu, D. Pál, and M. Sotáková. Learning low-density separators. In D. van Dyk and M. Welling, editors, *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR Workshop and Conference Proceedings, pages 25–32, 2009.
- K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft. When is nearest neighbor meaningful? In *Int. Conf. on Database Theory*, pages 217–235. Springer, 1999.
- D. Boley. Principal direction divisive partitioning. *Data Mining and Knowledge Discovery*, 2(4): 325–344, 1998.
- K. Fan. On a theorem of weyl concerning eigenvalues of linear transformations i. *Proceedings of the National Academy of Sciences of the United States of America*, 35(11):652, 1949.
- P. Fränti and O. Virtajoki. Iterative shrinking method for clustering problems. *Pattern Recognition*, 39(5):761–775, 2006.
- L. Hagen and A. B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE transactions on Computer-aided design of integrated circuits and systems*, 11(9):1074–1085, 1992.

- J. A. Hartigan. *Clustering algorithms*. John Wiley & Sons, Inc., 1975.
- J. A. Hartigan and P. M. Hartigan. The dip test of unimodality. *The Annals of Statistics*, 13(1): 70–84, 1985.
- P. J. Huber. Projection pursuit. *The annals of Statistics*, pages 435–475, 1985.
- A. Krause and V. Liebscher. Multimodal projection pursuit using the dip statistic. *Preprint-Reihe Mathematik*, 13, 2005.
- H.-P. Kriegel, P. Kröger, and A. Zimek. Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering. *ACM Transactions on Knowledge Discovery from Data*, 3(1):1–58, 2009.
- F. Leisch. A toolbox for k-centroids cluster analysis. *Computational Statistics & Data Analysis*, 51(2):526–544, 2006.
- A. S. Lewis and M. L. Overton. Eigenvalue optimization. *Acta numerica*, 5:149–190, 1996.
- J. R. Magnus. On differentiating eigenvalues and eigenvectors. *Econometric Theory*, 1(02):179–191, 1985.
- H. Narayanan, M. Belkin, and P. Niyogi. On the relation between low density separation, spectral clustering and graph cuts. In *Advances in Neural Information Processing Systems*, pages 1025–1032, 2006.
- D. Niu, J. G. Dy, and M. I. Jordan. Dimensionality reduction for spectral clustering. In *International Conference on Artificial Intelligence and Statistics*, pages 552–560, 2011.
- M. L. Overton and R. S. Womersley. Optimality conditions and duality theory for minimizing sums of the largest eigenvalues of symmetric matrices. *Mathematical Programming*, 62(1-3):321–357, 1993.
- N. Pavlidis, D. Hofmeyr, and S. Tasoulis. Minimum density hyperplanes. *arXiv preprint arXiv:1507.04201v2*, 2016.
- E. Polak. On the mathematical foundations of nondifferentiable optimization in engineering design. *SIAM Review*, 29(1):21–89, March 1987.
- A. Rosenberg and J. Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *EMNLP-CoNLL*, volume 7, pages 410–420. 2007.
- J. Schur. Bemerkungen zur theorie der beschränkten bilinearformen mit unendlich vielen veränderlichen. *Journal für die reine und Angewandte Mathematik*, 140:1–28, 1911.
- J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
- M. Steinbach, L. Ertöz, and V. Kumar. The challenges of clustering high dimensional data. In *New Directions in Statistical Physics*, pages 273–309. Springer, 2004.
- T. Tao and V. Vu. Random matrices have simple spectrum. *arXiv preprint arXiv:1412.1438*, 2014.
- S. K. Tasoulis, D. K. Tasoulis, and V. P. Plagianakos. Enhancing principal direction divisive clustering. *Pattern Recognition*, 43(10):3391–3411, 2010.
- S. Tong and D. Koller. Restricted bayes optimal classifiers. In *AAAI/IAAI*, pages 658–664, 2000.
- V. N. Vapnik and S. Kotz. *Estimation of dependences based on empirical data*, volume 40. Springer-verlag New York, 1982.

- U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.
- D. Wagner and F. Wagner. *Between min cut and graph bisection*. Springer, 1993.
- H. Weyl. Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, 1912.
- P. Wolfe. On the convergence of gradient methods under constraint. *IBM Journal of Research and Development*, 16(4):407–411, 1972.
- L. Xu, J. Neufeld, B. Larson, and D. Schuurmans. Maximum margin clustering. In *Advances in neural information processing systems*, pages 1537–1544, 2004.
- D. Yan, L. Huang, and M. I. Jordan. Fast approximate spectral clustering. In *Proceedings of the 15th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, pages 907–916. ACM, 2009.
- Q. Ye. Relative perturbation bounds for eigenvalues of symmetric positive definite diagonally dominant matrices. *SIAM Journal on Matrix Analysis and Applications*, 31(1):11–17, 2009.
- L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *Advances in neural information processing systems*, pages 1601–1608, 2004.
- K. Zhang, I. W. Tsang, and J. T. Kwok. Maximum margin clustering made practical. *IEEE Transactions on Neural Networks*, 20(4):583–596, 2009.
- T. Zhang, R. Ramakrishnan, and M. Livny. Birch: an efficient data clustering method for very large databases. In *ACM SIGMOD Record*, volume 25, pages 103–114. ACM, 1996.
- Y. Zhao and G. Karypis. Empirical and theoretical comparisons of selected criterion functions for document clustering. *Machine Learning*, 55(3):311–331, 2004.